REPORT DOCUMENTATION PAGE

1. AGENCY USE ONLY (Leave Blank)	2. REPORT DATE -6/21/99	3. REPORT TYPE AND DATES COVERED Final, 3/98-5/99

4. TITLE AND SUBTITLE
Statistical Signal Processing for Demining: Experimental Validation

5. FUNDING NUMBERS
DAAG55-98-1-0340

6. AUTHOR(S)
Leslie M. Collins

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Duke University, Box 90291, Durham, NC 2778-0291

8. PERFORMING ORGANIZATION REPORT NUMBER
2

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

U. S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

ARO 37953.1-EL

11. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not be construe an official Department of the Army position, policy or decision, unless so designated by the documentation.

12 a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

12 b. DISTRIBUTION CODE

13. ABSTRACT (Maximum 200 words)

Under the support provided by ARO in the form of a MURI for Humanitiarian demining, successful techniques for discriminating between mines and anthropic clutter have been developed using a statistical signal processing approac The improved performance provided by these algorithms has been validated using data obtained by DARPA. In ord determine whether these algorithms have wider application than the relatively high-metallic content mines used in th DARPA experiment, the Joint UXO Coordination Office (JUXOCO) was interested in augmenting the work begun un the MURI. JUXOCO is sponsoring a series of experiments designed to establish a performance baseline for metallic detectors. This baseline will be used to measure the potential improvements in performance offered by advanced sig processing algorithms. The goal of the work funded under this grant was to collect data from low-metal content mine using Geophex's GEM-3 sensor and to begin the development of improved detection algorithms. This report provide summary of the results obtained during the course of this study and a summary of experimental data acquisition methods.

14. SUBJECT TERMS

15. NUMBER OF PAGES
52

16. PRICE CODE

17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTR UL

19991101 112

Final Report

"Statistical Signal Processing for Demining: Experimental Validation

Leslie M. Collins

Duke University

Abstract

Under the support provided by ARO in the form of a MURI for Humanitiarian demining, successful techniques for discriminating between mines and anthropic clutter have been developed using a statistical signal processing approach. The improved performance provided by these algorithms has been validated using data obtained by DARPA. In order to determine whether these algorithms have wider application than the relatively high-metallic content mines used in the DARPA experiment, the Night Vision Laboratory (NVL) was interested in augmenting the work begun under the MURI. The Joint UXO Coordination Office at Ft. Belvoir, VA is sponsoring a series of experiments designed to establish a performance baseline for metallic mine detectors. This baseline will be used to measure the potential improvements in performance offered by advanced signal processing algorithms. The goal of the work funded under this grant was to collect data from low-metal content mines using Geophex's GEM-3 sensor and to begin the development of improved detection algorithms. This report provides a summary of the results obtained during the course of this study, a summary of experimental data acquisition methods, and the report generated for an early experiment conducted in conjunction with this project.

List of Appendices, Illustrations, and Tables

Background

Historically, electromagnetic induction (EMI) sensors have been used extensively to locate buried land mines. EMI sensors detect the metal that is present in such mines. However, land mines vary in their construction from metal-cased varieties with a large mass of metal to plastic-cased varieties with minute amounts of metal. Unfortunately, there is often a significant amount of metallic debris (clutter) present in the environment. Consequently, EMI sensors that utilize traditional detection algorithms based solely on the metal content suffer from large false alarm rates. When the sensor sensitivity is set to detect the very small amount of metal found in typical plastic-case mines, these false alarm rates become unacceptably high. In addition, for high-metallic mines, it is commonly assumed that the effects of the soil are negligible with respect to the sensor response. To date, no testing has been done to validate the assumption that the effects of soil are negligible when detecting low-metallic content mines.

Under the support provided by ARO in the form of a MURI for Humanitiarian demining, successful techniques for discriminating between mines and anthropic clutter have been developed using a statistical signal processing approach [1-4]. The improved

performance provided by these algorithms has been validated using data obtained by DARPA [5]. In order to determine whether these algorithms have wider application than the target set used in the DARPA experiment, the Night Vision Laboratory (NVL) was interested in augmenting the work begun under the MURI. The Joint UXO Coordination Office at Ft. Belvoir, VA is sponsoring a series of experiments designed to establish a performance baseline for metallic mine detectors. This baseline will be used to measure the potential improvements in performance offered by advanced signal processing algorithms.

Problem Studied

The long-term goal of this research is to determine whether enhanced signal processing algorithms derived using signal detection theory (SDT) applied to data collected with EMI sensors (*e.g.*, Geophex's GEM-3 sensor and the Schiebel PSS-12 sensor) substantially reduce false alarm rates over traditional approaches. To achieve this goal, we performed several preliminary experiments and analyses to show proof of concept. These included

(1) performing an experiment at Geophex, Ltd. to determine whether the GEM-3 could be used to measure signatures from buried low-metal mines,

(2) gathering signature data at a calibration site located at Fort A. P. Hill,

(3) comparing the false alarm rate of the SDT-based algorithms to energy-based algorithms at the same detection rate on data collected at a calibration site at Fort A. P. Hill, and

(4) gathering signature data at the all data points in the blind test grid located at Fort A. P. Hill.

NVL was responsible for creating and maintaining the test site at Fort A. P. Hill. The targets that were emplaced included both low-metal and high-metal anti-personnel mines. The targets were emplaced on an x-y grid with approximately 1,000 intersections, with the intersections spaced adequately to avoid target interference. Approximately 100 of the intersections contained targets. A calibration area was also made available adjacent to the test site which contained a few targets for initial system calibration.

This report includes a description of

(1) a summary of the results obtained during the course of this study

(2) a summary of experimental data acquisition methods, and

(3) the report generated for the Geophex experiment (in Appendix A).

Also associated with this work are

(1) "Raw" sensor output data from collection efforts on ZIP disks

(2) An excel file listing the data collection parameters for each grid point measured.

These items have been delivered to JUXOCO for dissemination on their web site.

Summary of Results

1. Geophex Experiment

The main goal for this preliminary experiment was to use the GEM-3 sensor to collect digital, multi-frequency signal data from mines in free space and mines buried in soil. Extensive measurements of the background were taken, both in free space and in soil, to determine differences between (a) the background with mines present, and (b) the background without mines present. This data also allows us to examine the stability of the sensor response. Geophex recently obtained data with the GEM-3 that indicate that, in free space, it can both detect and discriminate various low-metal content mines. We wanted not only to verify Geophex's free space data with the same mine types but also to determine the effect of burial in soil on these mine signatures. We also included two high-metallic content mines in our target set. These large metal mines were used to confirm a commonly held assumption: for "large" metal objects the effects of burial in soil are negligible on the sensor response. By comparing the signatures of these large metal mines in free space and when buried, we could evaluate the validity of this assumption. In addition, we could also determine whether the signatures of low-metallic content mines are the same as in free space as when they are shallowly buried. To date, the validity of this assumption has not been tested for low-metallic content mines.

A description of the GEM-3, the experimental protocol, and a detailed analysis of the results of this experiment are contained in the report provided in Appendix A. A summary of the results are:

(1) The presence of soil does affect the signature recorded by the GEM-3 for low-metallic content mines, and does not affect the signature for high-metallic content mines.
(2) It is possible to detect low-metal land mines using the GEM-3.
(3) The measurements obtained are statistical in nature, not deterministic. Detectors, or algorithms, which effectively incorporate the stochastic character of the signals should be able to out-perform traditional detection algorithms.
(4) The sensor experiences some drift in its response. Drift must be considered both in gathering data and in the analysis. At a minimum, background measurements must be made during data collection for accurate background correction.

2. Calibration area data collection

Hand-held mine detector experiments were conducted between June and July 1998 at Range 71A of Fort A. P. Hill. The test site and preparation are described in detail in [7]. A calibration area was established consisting of known targets and known, well characterized, clutter emplaced at known locations. The calibration area consisted of 5

lanes, labeled A through E, each 25 meters long and 1 meter wide. This area was located adjacent to the main test grid, and contained at least one of each mine contained in the main test grid. Data collection in the calibration area was performed by the PI and a graduate student, Ms. Ping Gao. Mr. John Moulton and Mr. John Carey of EOIR assisted with all aspects of the data collection.

1 meter square data collection templates were constructed for placement at each grid location in both the calibration lane and the main test grid to facilitate signature collection. The template was constructed of foam, and had a series of marks for use in locating the sensor head at the appropriate positions. 10 spatial positions were measured in each grid point. For grid points in which the GEM-3 response did not overload, these points were located at –4", -2", 0", 2", and 4" from the center in both a vertical and a horizontal orientation [7]. Collection points were labeled 1 – 5 from top to bottom and 6 – 10 from left to right. When the sensor response overloaded, the sensor head was raised 24" above the ground and the spatial samples were taken every 4" as opposed to every 2". Signatures were measured at 20 frequencies spaced logarithmically between 270 Hz and 23,790 Hz. Each signature was measured twice. All measurements were recorded in a log-file in ascii format (see section 5 below for more detail). The log file was named according to the grid. For example, the data collected for grid square B2 was saved in a file called 'B2.ppm'.

Data was collected in the calibration area four times: on June 9-11, July 6-7, July 27, and September 30, respectively. Data was measured at each grid point where either a target or clutter had been emplaced. Following the collection of signatures at each of the 10 spatially distributed points corresponding to one grid square, a background measurement was taken at a blank grid point. These "background" measurements were used to track the sensor drift described above, and in the report attached in Appendix A. The background measured after a grid square is labeled according to the grid square it follows. For example, the background measured following the collection in square B2 would be called 'bg_b2.ppm'.

Following data collection, the data were hand-corrected for any data drop-outs that occurred as a result of the serial connection between the GEM-3 and the PC controller. These data were then analyzed as described in the following section.

3. Performance enhancements using calibration data

A series of simple detectors were designed to analyze potential false alarm rate improvements using the data obtained in the calibration lanes. Prior to implementing the detectors on the data measured at each grid point, the background must be removed. An example of the variability of the background is shown in Figure 1. Figure 1 illustrates the variability in the in-phase component of the background response measured on 7/7/98. Responses are plotted as PPM versus log-frequency – see [7] for a description of the PPM unit. Clearly, substantial variation in the background signature occurs over the course of the day, especially in the high frequencies.

Figure 1: Variation in in-phase component of the GEM-3 signature on 7/7/98

Several approaches for background correction were analyzed. In the first, the background measured immediately prior to the measurement of a grid point was subtracted from each spatial signature. In the second, the background measured immediately after the measurement of a grid point was subtracted from each spatial signature. The third approach used the average of the background signatures measured immediately before and after measuring the signatures in the grid. The final approach utilized a linear prediction of the background obtained from the background measured before and the background measured after the data collection for a particular the grid square in order to provide background correction.

The performance of the background correction algorithms was analyzed by performing one measurement in which a background was taken between every spatial measurement, *i.e.* the order of measurements was Background 1 – Spatial position 1 – Background 2 - Spatial position 2 - ... - Background 10 – Spatial position 10 – Background 11. Using this approach the true background for a particular spatial position could be very accurately estimated from the bordering background measurements. The "actual" background-corrected signatures could then be compared to those obtained using the four methods, corresponding to using Background 1, Background 11, the average of Background 1 and 11, and linear prediction using Backgrounds 1 and 11 respectively.

As expected, using only Background 1 produced substantial errors in the spatial positions measured later (7 —10). Similarly, using only Background 11 produced substantial errors in the spatial positions measured earlier (1 – 3). Utilization of the average background produced fair estimates in the middle spatial positions (4-6), but poorer estimates in the

earlier and later positions. The linear prediction method provided the best approach to background correction. Examples of the background corrected estimates for the signatures are shown for one grid point (K-32) at the first, fifth, and tenth spatial position respectively in Figures 2-4. The true signature is shown with the solid line, the dashed line illustrates the prediction using the previous background, the dashed-dotted line illustrates the prediction using the following background. The black line shows the signature obtained using the linear prediction. Both in-phase (red) and quadrature (blue) data are shown. Clearly, the prediction method works best, especially at the higher frequencies.

**Figure 2. Background correction example for grid point K32 at the first spatial position. "Previous" background approach is coincident with true solution.**

Figure 3. Background correction example for grid point K32 at the fifth spatial position.


Figure 4. Background correction example for grid point K32 at the tenth spatial position. "Following" background approach is coincident with true solution.

Once background correction has been adequately achieved, various algorithms can be applied to the data. In the main grid, grid points can consist of targets, clutter, or "blanks", i.e., empty spots. In the calibration lane, only grid points consisting of targets and clutter were measured, so future work will included simulating blanks for analysis with the calibration lane data.

We consider two classes of algorithms, those which are based solely on the energy measured at a particular point and those which process the measured signature. We further divide these two classes into algorithms which operate only on the data measured at the center of the grid, and those measured at the various spatial points. In this way, we can quantify the performance of a "baseline", which we consider to be center-point energy-based algorithms. We can also quantify performance gains associated with using the entire frequency-domain signature, as well as incorporation of spatial information. Finally, when considering energy-based algorithms, we subdivide those into various ad-hoc procedures versus statistically-based Bayesian procedures. When considering signature-based algorithms, we consider both statistically-based Bayesian procedures as well as model-based Bayesian procedures. This classification system is illustrated in Figure 5. White blocks have been pursued with the calibration data set, gray blocks remain future work.
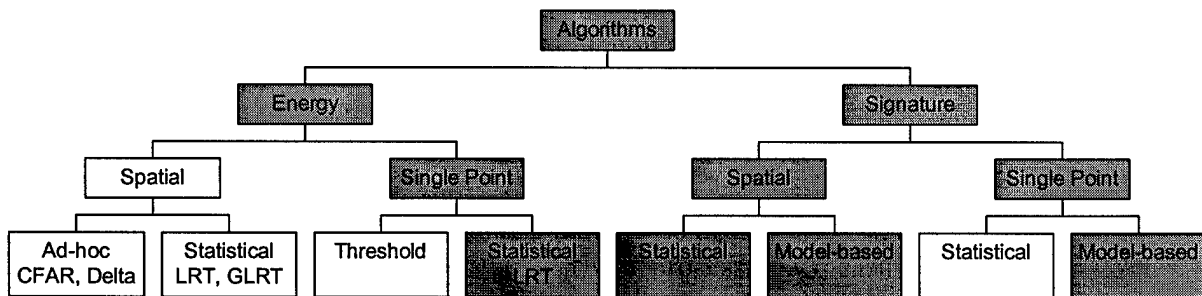
Figure 5. Algorithm classification chart. White boxes have been implemented in the calibration area, gray boxes remain future work.

Figures 6 and 7 illustrate ROC curves for various energy-based processors. Figure 6 illustrates the performance of a simple energy threshold for the center point (red line), a CFAR processor which thresholds the difference between the center point and the average of the adjacent points (green line) [8], the delta processor, which thresholds the number of times the center point was larger than the surrounding eight points (magenta line) [9], and a GLRT/matched filter (blue line) [1]. Clearly, baseline performance follows the chance diagonal, and the GLRT provides some performance improvement. The GLRT shown in Figure 6 assumes Gaussian distributed data. The 2D-GLRT described in [1] estimates the statistics of the data and derives a GLRT based on the spatial statistics. Figure 7 illustrates the performance of the 2D-GLRT (blue line) as compared to a thresholded energy (red line). Clearly, it is possible to improve performance above the baseline using these techniques.

Figure 8 illustrates the ROC for a statistically-based Bayesian processor operating on the signature measured at the center point (blue line). The processor is trained on the

signature measured at spatial position 3 and tested on the signature measured at spatial position 8. Performance is compared to the baseline energy detector (red line). Clearly, incorporating the frequency-domain signature information shows potential for substantially improving detector performance.

Ad-hoc spatial processors

Figure 6. Comparison of the performance of various energy-based processors.

Bayesian processor using signature at point 8

Figure 7. Comparison of the performance of a Bayesian spatial processor using energy data to the baseline energy detector performance.

**Bayesian spatial processor (2D-GLRT) using energy only**

Figure 8. Comparison of the baseline performance (energy detector) to two processors which incorporate the physical nature of the signal into the processing.

4. Main grid data collection

Hand-held mine detector experiments were conducted in the main grid between July and October 1998 at Range 71A of Fort A. P. Hill. The test site and preparation are described in detail in [7], and the data collection mechanisms are described above. The main test grid was established consisting of approximately 100 targets, and an unknown number of clutter objects and blank squares. The main grid consisted of 20 lanes, labeled A – T, each 49 meters long and 1 meter wide. Data collection in the calibration area was primarily performed by the PI and a graduate student, Ms. Ping Gao. Two other Duke University graduate students, Mr. David Ferguson and Ms. Yingyi Tan, also assisted in the data collection at various times. Mr. John Moulton and Mr. John Carey of EOIR assisted with all aspects of the data collection. The format of the data is described briefly below. At each grid point, the data and time of the collection were recorded. Data was collected 7/27-31/98, 8/10-14/98, 9/21-25/98, 9/28-30/98, and 10/27-29/98. Background measurements were taken at one of four grid squares: E11, E39, O11 and O39. These squares were set aside as known "blanks" by JUXOCO. The closest blank was used as the background for each grid square.

The calibration data and the main grid data have been transferred to JUXOCO to be placed on their web site. All of the data from each of the grid points and background measurements were supplied. In addition, an excel file which lists the grid point, and the associated date of collection, time of collection, location of associated background measurement, and file names for the grid point data file, previous and following background files are listed.

5. Data formats

The GEM-3 can be programmed to sequentially record the magnetic field at a set of user defined frequencies. To select these frequencies, the operator selects

(1) a minimum and maximum frequency,
(2) the number of frequencies, and
(3) linear or logarithmic spacing of frequencies.

Based on these parameters, the frequencies are generated automatically by the GEM-3 operating system. To gather the data at Fort A. P. Hill, we selected 20 logarithmically spaced frequencies. When the induced magnetic field is recorded, both in-phase and quadrature (measured in ppm) are measured. The GEM-3 can also be programmed to automatically repeat the measurement an arbitrary number of times. At Fort A. P. Hill, we repeated each measurement 2 times. Thus, at a particular spatial location, 2 measurements were taken. Therefore, for an entire grid square, 20 measurements were taken.

In this experiment, we operated the GEM-3 remotely. We used a laptop PC to transmit data acquisition parameters to the GEM-3 and to record the measured responses through the serial port. Prior to data acquisition, we allowed the sensor to warm up for at least ten minutes. As the data is acquired, it is written to a log file on the hard disk. The log file consists of a header that lists the acquisition parameters, followed by the recorded data. Each line of the recorded data lists the acquisition number, measurement number (*i.e..*, 1 to 2), frequency, in-phase response, quadrature response, and an error flag. Commas separate the recorded data. Note, the first number, called 'acquisition number' is a number internally maintained by the GEM-3. When the sensor is turned on, it is initialized to 1. Every time a scan is completed, it is incremented by the GEM-3 operating system. This number is not necessarily indicative of the order of the scans, since the number is incremented even if a scan is aborted, which sometimes occurs if errors are detected.

Table I shows ans example of the recorded data. A portion of the second measurement is also shown. The measurements are listed as acquisition number 5. As shown, additional measurements are appended to the end of the listing of the first measurement. The acquisition number does not change, the measurement number is incremented, the list of frequencies does not change, and each in-phase and quadrature sample is recorded along with the error flag.

Acquisition Number	Measurement Number	Frequency	In-phase Data	Quadrature Data	Error Flag
5	1	270	48	-79	0
5	1	330	198	-55	0
5	1	420	180	49	0
5	1	540	205	85	0
5	1	690	315	94	0
5	1	870	445	101	0
5	1	1110	972	83	0
5	1	1410	1369	35	0
5	1	1800	2069	-48	0
5	1	2250	2939	-340	0
5	1	2850	4253	-827	0
5	1	3630	6058	-1727	0
5	1	4590	8387	-3268	0
5	1	5820	14949	-9999	0
5	1	7350	18874	-16281	0
5	1	9330	25686	-36616	0
5	1	11790	27499	-51767	0
5	1	14940	27463	-70573	0
5	1	18930	24938	-94418	0
5	1	23970	9202	-162973	0
5	2	270	49	-75	0
...	...	...	...	...	...

Table I: An example of a raw data sample recorded by the GEM-3.

The GEM-3 records in-phase and quadrature data as a function of frequency in ppm. These data are based on a current measurement, as opposed to a magnetic field measurement. Therefore, to convert these measurements to data that are proportional to magnetic field strength, an inverse dependence on the frequency and a 90-degree phase shift must be incorporated (see Appendix I for details). Thus, the following calculations should be used for each of the measurements taken:

$$I(f) = -\frac{Q_s(f)}{f}$$
$$Q(f) = \frac{I_s(f)}{f}$$

Here, f is frequency, $I_s(f)$ and $Q_s(f)$ are the in-phase and quadrature samples reported by the GEM-3, and $I(f)$ and $Q(f)$ are the calculated in-phase and quadrature samples that are proportional to the magnetic field.

List of all publications and technical reports

(1) Gao, P., Collins, L. M., Garber, P., Geng, N. and Carin, L., "Classification of landmine-like metal targets using wideband electromagnetic induction", submitted to IEEE Trans. Geoscience and Remote Sensing, November, 1998.
(2) Collins, L. M. and Gao, P., "Progress Report I: Signatures of Land Mines in Soil and in Air: Are They Different? What are the Characteristics of the Sensor Noise?", August 5, 1998.
(3) Gao, P., Collins, L., Weaver, D., Keiswetter, D., and Won, I.J., "Enhanced Detection of Landmines using Broadband EMI Data", submitted to Detection and Remediation Technologies for Mines and Minelike Targets IV Conference, 1999 International Symposium on Aerospace/Defense Sensing and Controls.
(4) Gao, P., "Improved Signal Processing Approaches to Landmine Detection", Dissertation Research Proposal, Department of Electrical and Computer Engineering, Duke University, December 4, 1998.

List of all participating personnel

(1) Leslie M. Collins
(2) Ping Gao

Bibliography

[1].   Collins, L., P. Gao, and L. Carin, *An Improved Bayesian Decision Theoretic Approach for Land Mine Detection*. IEEE Trans. Geosc. Remote Sens., 1997 (submitted).

[2].   Gao, P. and L. Collins, *A 2-Dimensional Generalized Likelihood Ratio Test for Land Mine Detection.* Signal Processing Letters, 1998 (submitted).

[3].   Gao, P. and L. Collins. *Improved signal processing approaches for landmine detection.* in *SPIE.* 1998. Orlando, FL.

[4].   Geng, N., et al., *Wideband Electromagnetic Induction for Metal-Target Identification: Theory, Measurement and Signal Processing.* IEEE Trans. Geoscience and Remote Sensing, 1998(submitted).

[5].   George, V., et al., *Background Data Collection Plan,* . 1996, DARPA/Defense Science Office.

[6].   Won, I.J., D.A. Keiswetter, and D.R. Hansen, *GEM-3: A Monostatic Broadband Electromagnetic Induction Sensor.* J. Envir. Engin. Geophysics, 1997. **2**: p. 53-64.

[7].   *Handheld Metallic Mine Detector Performance Baselining Collection Plan,* JUXOCO, Fort Belvoir, VA, May, 1998.

[8].   Holmes, Q. A., Schwartz, C. R., and Seldin, J. H., *"Adaptive Multispectral CFAR Detection of Landmines"*, Proceedings of the SPIE, vol. 2496, Orlando, FL, 1995.

[9].   Gelenbe, E. and Kocak, T., *"Area-based Results for Mine Detection"*, submitted to IEEE Trans. Geoscience Remote Sensing, 1997.

Appendix A – Geophex Experiment Report

Statistical Signal Processing for Demining: Experimental Validation

Progress Report I: Signatures of Land Mines in Soil and in Air: Are They Different? What are the Characteristics of the Sensor Noise?

An analysis of data gathered with the GEM-3 at Geophex, Ltd. from six different mines

Leslie Collins and Ping Gao
Electrical and Computer Engineering Department
Duke University

August 5, 1998

Prepared for:

Mr. Dick Weaver, Deputy Director
Joint Unexploded Ordnance Center of Excellence

I. Introduction

A. Background

Historically, electromagnetic induction (EMI) sensors have been used extensively to locate buried land mines. EMI sensors detect the metal that is present in such mines. However, land mines vary in their construction from metal-cased varieties with a large mass of metal to plastic-cased varieties with minute amounts of metal. Unfortunately, there is often a significant amount of metallic debris (clutter) present in the environment. Consequently, EMI sensors that utilize traditional detection algorithms based solely on the metal content suffer from large false alarm rates. When the sensor sensitivity is set to detect the very small amount of metal found in typical plastic-case mines, these false alarm rates become unacceptably high. In addition, for high-metallic mines, it is commonly assumed that the effects of the soil are negligible with respect to the sensor response. To date, no testing has been done to validate the assumption that the effects of soil are negligible when detecting low-metallic content mines.

The Joint UXO Coordination Office at Ft. Belvoir, VA is sponsoring a series of experiments designed to establish a performance baseline for metallic mine detectors. This baseline will be used to measure the potential improvements in performance offered by advanced signal processing algorithms. This report describes a preliminary experiment that is part of the larger series of experiments to develop performance baselines for mine detection by EMI sensors.

B. Approach

Our approach to lowering false alarm rates is to develop advanced algorithms that take advantage of additional information present in the received signal from EMI sensors. Conventional algorithms have not exploited this information. The general concept is to separate the mine's return signal from both the background and the clutter signal returns. The signal that is measured by an EMI system consists of either (1) a "background" return alone (due to soil, etc.), (2) the superposition of a background return and a return from a mine, or (3) the superposition of a background return and a return from clutter. Thus, if the background signal is subtracted from the measured signal, we can "see" the mine or clutter signature and potentially discriminate between the two.

To implement the advanced algorithms, we will use signal detection theory. Signal detection theory affords a powerful tool for designing algorithms to both detect and discriminate signals of interest in the presence of noise and man-made, or anthropic, clutter. To utilize this tool, it is necessary to define (1) a set of hypotheses to be tested and (2) the statistical nature of the data that is associated with each of the hypotheses. In this preliminary experiment, we ignore the possibility of clutter and assume that the signals that are measured correspond to the response of the sensor either to a mine in a background environment or to the background environment alone. Future experiments will address the clutter issue in detail.

Thus, for this preliminary experiment, there are only two hypotheses: H_1: "mine present in background" and H_0: "background alone". It is assumed that the response of the sensor under either of the hypotheses is subject to some amount of "noise". This noise can be ascribed either to the electronics of the sensor itself, or to random variations in environmental parameters. In either case, it is important to characterize the statistics of the noise in order to formulate the detection algorithm prescribed by signal detection theory. The first step in this process is to characterize the probability density function (pdf) that describes the noise process (*e.g.,* Gaussian). This characterization allows us to obtain a general description of the statistical

parameters of the noise and its pdf, such as its mean and variance. We also generate a pdf for the target responses in different backgrounds. Next, it is important to determine whether or not these statistics are stable, or stationary, *i.e.*, whether the statistics vary as a function of the signal parameters themselves. For this preliminary experiment, we examined this stability assumption over the following parameters:

(1) the absolute amplitude level of the received signal, which may correspond to whether or not a mine is present

(2) the amount of time the sensor has been operating, corresponding to something akin to "drift" in the mean response of the sensor, and

(3) the frequency at which a measurement is made.

II.   Experiment Description

A. Goals

The main goal for this preliminary experiment was to use the GEM-3 sensor to collect digital, multi-frequency signal data from mines in free space and mines buried in soil. Extensive measurements of the background were taken, both in free space and in soil, to determine differences between (a) the background with mines present, and (b) the background without mines present. This data also allows us to examine the stability of the sensor response, as discussed above. Geophex recently obtained data with the GEM-3 that indicate that, in free space, it can both detect and discriminate various low-metal content mines. We wanted not only to verify Geophex's free space data with the same mine types but also to determine the effect of burial in soil on these mine signatures. We also included two high-metallic content mines in our target set. These large metal mines were used to confirm a commonly held assumption: for "large" metal objects the effects of burial in soil are negligible on the sensor response. By comparing the signatures of these large metal mines in free space and when buried, we could evaluate the validity of this assumption. In addition, we could also determine whether the signatures of low-metallic content mines are the same as in free space as when they are shallowly buried. To date, the validity of this assumption has not been tested for low-metallic content mines.

B. Mine Descriptions

To address the goals of the experiment, six inert land mines were provided by NVESD through the Joint UXO Coordination Office. These were actual mines that were "down-loaded" (all explosives removed). For the low-metal content mines, care was taken to insure that all the correct metal parts were present in the proper location and orientation. Table I lists the mines that were tested, along with their relevant characteristics [1]. Both high-metal and low-metal content mines were evaluated in order to measure the effects of the soil and the statistics of the noise process as a function of metal content.

Mine Name	Type	Diameter (cm)	Description and Metal Content
Valmara	AP, M	10	Italian bounding fragmentation mine. Plastic case over large metal canister. Total metal: 2800 g
VS50	AP, LM	9	Italian round plastic-cased blast mine. Mine case empty of any fill. Total metal: 18.21 g

M14	AP, LM	5.6	US and Indian manufactured plastic bodied blast mine. Case empty of any fill. Total metal: 0.6 g
M19	AT, LM	33	American rectangular plastic blast mine. Test mine case filled with room temperature vulcanized material (RTV). Total metal: 0.94g
VS2.2	AT, LM	23	Italian plastic blast mine. Case empty of any fill. Total metal: 3.29 g
TS50	AP, LM	9	Italian plastic-cased cylindrical blast mine. Case empty of any fill. Total metal: 4.41 g

Table I. Mines tested in the experiment. AP = Antipersonnel mine, AT = Antitank mine. M = High metal content, LM = Low metal content.

C. Sensor

Physical nature of the sensor response. The GEM-3 is a prototype wide-band frequency-domain EMI sensor manufactured by Geophex, Ltd. The GEM-3 uses a pair of concentric, circular coils to transmit a continuous, wideband, digital electromagnetic waveform [2]. The resulting field induces a secondary current in the earth as well as in any buried objects. The set of two transmitter coils has been designed so that they create a zone of magnetic cavity at the center of the two coils. A third receiving coil is placed within the magnetic cavity so that it senses only the weak secondary field returned from the earth and buried objects.


Figure 1. Model of the GEM-3 receive coil. Red path indicates the path of the induced current and thus the line contour for integration; blue shaded area indicates surface area for integration. Entire wire diameter is noted.

The receiving coil can be modeled as shown in Figure 1. Faraday's Law gives the equation governing response induced in the receiving coil:

$$emf = \oint_c E \cdot dl = -\frac{d}{dt}\iint_s B \cdot \hat{n}ds = -\frac{d\Phi}{dt}$$

where *emf* is the electromotive force, E is the electric field intensity, l is the path length through the wire, B is the magnetic flux density, $\hat{n}$ is the unit vector perpendicular to ds, ds is the

4

differential surface element on the surface S circumscribed by the coil, and Φ is the magnetic flux. Since the electric field, E, is zero inside the wire, the above equation can be rewritten as

$$V_b - V_a = \frac{d}{dt} \oiint_S B \cdot \hat{n} ds = \frac{d\Phi}{dt}.$$

Using Fourier theory, this derivative relationship in the frequency domain is transformed into a linear relationship in the frequency domain as

$$V_b - V_a = j\omega\Phi(\omega).$$

Thus, the measured frequency-domain voltage across the loop is proportional to the magnetic flux and, correspondingly, to the magnetic flux density.

In practice, the version of the GEM-3 used in this experiment measures the current through a resistor in series with the receiving loop. Measuring the current is equivalent to measuring the voltage across the loop. This current is 90 degrees out of phase with the voltage, and therefore the magnetic flux density. The in-phase and quadrature components at each frequency of interest are obtained by convolving (multiplying point by point and adding) the received time-series with a sine time-series (for in-phase) and cosine time-series (for quadrature) at the frequency of interest. The data obtained from the convolution are converted into units called parts-per-million (ppm) defined as

$$ppm = \frac{measured\_data}{calibration\_data} x10^6.$$

The calibration data at each frequency had previously been obtained (and programmed into the GEM-3) by hanging the GEM-3 "from the top of a tall tree" [2,3] and measuring the sensor output. This normalization attempts to remove the system impulse response from the measured data. The ppm data are logged by the sensor for each frequency as in-phase and quadrature components. The phase relationships and frequency dependence are not corrected at this point.

Operational Issues. The GEM-3 can be programmed to sequentially record the magnetic field at a set of user defined frequencies. To select these frequencies, the operator selects

(1) a minimum and maximum frequency,
(2) the number of frequencies, and
(3) linear or logarithmic spacing of frequencies.

Based on these parameters, the frequencies are generated automatically by the GEM-3 operating system. The operator has the option to make modifications to the list of frequencies. The process of recording the induced magnetic field (in-phase and quadrature measured in ppm) at each of the specified frequencies will be referred to as a *measurement* (see Figure 2). For example, if the GEM-3 is programmed to record data at 4 kHz and 10 kHz, then a measurement would consist of in-phase and quadrature data (in ppm) at each frequency, *i.e.,* four data values would be recorded. Each of these individual data values will be referred to as a *sample*. The GEM-3 can also be programmed to automatically repeat the measurement an arbitrary number of times. For example, the GEM-3 could be programmed to measure the response at 4 kHz and 10 kHz 5 times. Multiple repetitions of a measurement for a particular experimental setup will be referred to as a

scan. A scan can be used to assess the repeatability of a measurement, and thus the statistics of the noise process.

To summarize the terminology, assume that the GEM-3 has been programmed to record the magnetic field induced at 4 kHz and 10 kHz and to repeat the data collection 5 times. A *sample* is a particular response that is recorded by the GEM-3, *e.g.* the 4 kHz in-phase component. 5 *measurements* would be taken, each of which consists of 4 samples. The *scan* is the entire set of 5 measurements. This terminology is illustrated in Figure 2.
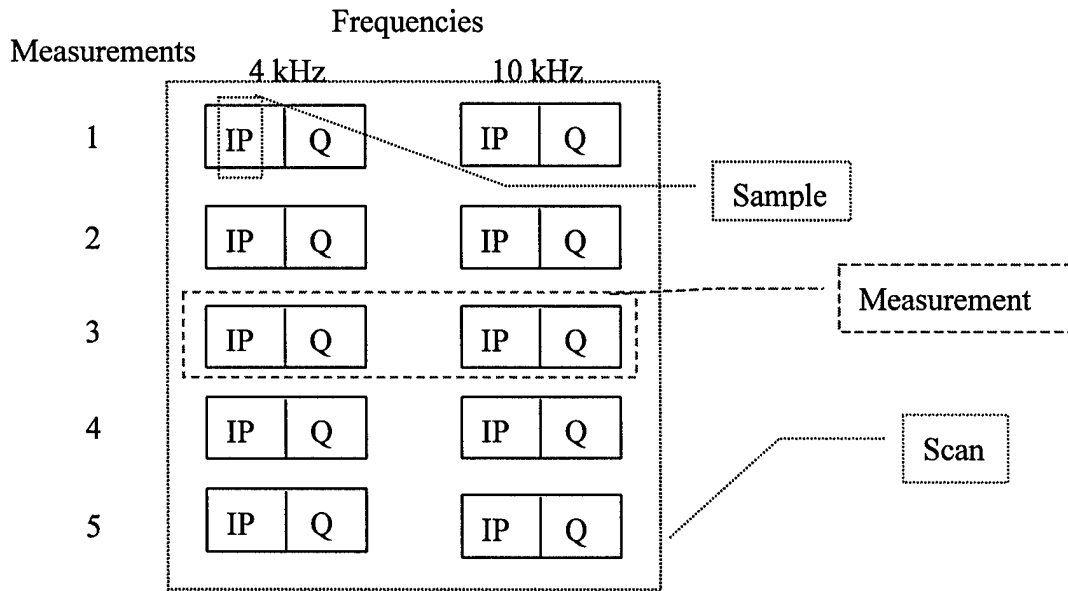

Figure 2. Illustration of terminology. IP denotes in-phase component, Q denotes quadrature component.

In a "noise-free" sensor, each of the samples taken above would be identical across all of the measurements. For example, the 4 kHz in-phase sample taken in measurement would be exactly the same as the 4 kHz in-phase sample taken in measurements 2-5. However, due to sensor noise, these samples will rarely be identical. Thus, by evaluating the statistics of a particular sample across the set of measurements, we can characterize the noise process. This can lead to an estimate of the pdf of the noise process, as discussed in Section III.

In this experiment, we operated the GEM-3 remotely. We used a laptop PC to transmit data acquisition parameters to the GEM-3 and to record the measured responses through the serial port. Prior to data acquisition, we allowed the sensor to warm up for at least ten minutes. As the data is acquired, it is written to a log file on the hard disk. The log file consists of a header that lists the acquisition parameters, followed by the recorded data. Each line of the recorded data lists the acquisition number, measurement number (*e.g.*, 1 to 5 in the example shown in Figure 2), frequency, in-phase response, quadrature response, and an error flag. Commas separate the recorded data. Note, we call the first number 'acquisition number' as opposed to scan number. It is a number internally maintained by the GEM-3. When the sensor is turned on, it is initialized to 1. Every time a scan is completed, it is incremented by the GEM-3 operating system. This number is not necessarily indicative of the order of the scans, since the number is incremented even if a scan is aborted, which sometimes occurs if errors are detected.

Table II shows an example of the first measurement of the M19 scan recorded in free-space. A portion of the second measurement is also shown. The measurements are listed as acquisition number 5, since the warm-up scan was number 1 (a background measurement), and the M14 was the 4th data set gathered after the warm-up (see following section for additional details regarding the order in which the scans were gathered). As shown, additional measurements are appended to the end of the listing of the first measurement. The acquisition number does not change, the measurement number is incremented, the list of frequencies does not change, and each in-phase and quadrature sample is recorded along with the error flag.

Acquisition Number	Measure ment Number	Frequency	In-phase Data	Quadrature Data	Error Flag
5	1	90	7	18	0
5	1	120	12	26	0
5	1	150	18	28	0
5	1	180	761	-284	0
5	1	240	48	51	0
5	1	300	198	-55	0
5	1	360	105	63	0
5	1	420	180	49	0
5	1	510	205	85	0
5	1	630	315	94	0
5	1	750	445	101	0
5	1	900	585	65	0
5	1	1110	972	83	0
5	1	1320	1369	35	0
5	1	1620	2069	-48	0
5	1	1950	2939	-340	0
5	1	2370	4253	-827	0
5	1	2880	6058	-1727	0
5	1	3480	8387	-3268	0
5	1	4230	11397	-5884	0
5	1	5130	14949	-9999	0
5	1	6240	18874	-16281	0
5	1	7560	22620	-24987	0
5	1	9150	25686	-36616	0
5	1	11100	27499	-51767	0
5	1	13440	27463	-70573	0
5	1	16320	24938	-94418	0
5	1	19770	19226	-124453	0
5	1	23970	9202	-162973	0
5	2	90	8	17	0
...	...	...	...	...	...

Table II: An example of a raw data sample recorded by the GEM-3.

D. Data Collection – General Protocol

The goals of this preliminary experiment were to assess the detectability of the mines in both free space and when buried, to determine whether the signal measured for each mine is identical in soil and free space, and to investigate the statistics of the noise process associated with the sensor. To assess the detectability of the mines and the effects of the ground on a mine's "signature", the response of the sensor to each mine as a function of frequency was determined both in free space and when buried in the soil. In the initial experiment, we characterized the frequency response in the range of 2 kHz to 24 kHz. 25 frequencies were linearly spaced in this range.

As was discussed previously, when the mine is present, the response of the sensor consists of a response both to the mine and a response to the background. To obtain the response due to the mine alone, it is necessary to determine the response of the sensor to the background alone. Therefore, in this experiment, we interspersed recordings (scans) of the response of the sensor to the background (no-mine) with recordings (scans) of the response of the sensor to the mine embedded in its background. Responses were initially recorded in free-space, then responses were recorded in the soil. To investigate the statistics of the noise process, including its stability, 50 measurements were taken in each scan. The experimental protocol can be summarized as follows:

(1)   Sensor warmed up for 10 minutes, and was never subsequently turned off.
(2)   Background scan: 50 measurements across the frequency range of 2,000 Hz to 24,000 Hz in linear steps. No mine present.
(3)   Mine scan: 50 measurements in the same frequency range. Sensor located in free space approximately 6 inches from the top of the mine.
(4)   Repeat steps 2-3 with each mine until all free space scans completed.
(5)   Final background scan (50 measurements, same frequency range) in free space.
(6)   Background scan: 50 measurements across the frequency range of 2,000 Hz to 24,000 Hz with no mine present with sensor located approximately 4 inches from the soil.
(7)   Bury mine, top of mine 1 inch deep. Tamp soil.
(8)   Mine scan: 50 measurements in the same frequency range. Sensor located approximately 4 inches from the soil.
(9)   Excavate mine, tamp soil.
(10)   Repeat steps 6 through 9 with each mine until all soil scans completed.
(11)   Final background scan (50 measurements, same frequency range) in soil.

Under these conditions, a complete scan (50 repeated measurements, each measurement taken at 25 frequencies) required approximately 10 minutes. The entire free space data collection was completed in approximately 3.5 hours, and the soil data collection required approximately 4.5 hours. The sensor was never turned off during the experiment.

A graphical illustration of the experiment flow is contained in Figure 3 for the case when 25 measurements are taken for each scan. This experimental protocol was implemented on April 29, 1998 and April 30, 1998, however 50 samples were taken for each scan. During the soil data collection phase, a concurrent experiment that was being performed at Geophex forced a temporary relocation of the sensor. Following this relocation, a substantial offset was observed in the data obtained during the background scans. Upon investigation, Geophex engineers determined that a screw that attaches the GEM-3 sensor head to its supports had become loose. As a result, the subsequent data that was obtained was considered to be invalid, and the experiment was repeated, with the modifications discussed below, on May 14, 1998.
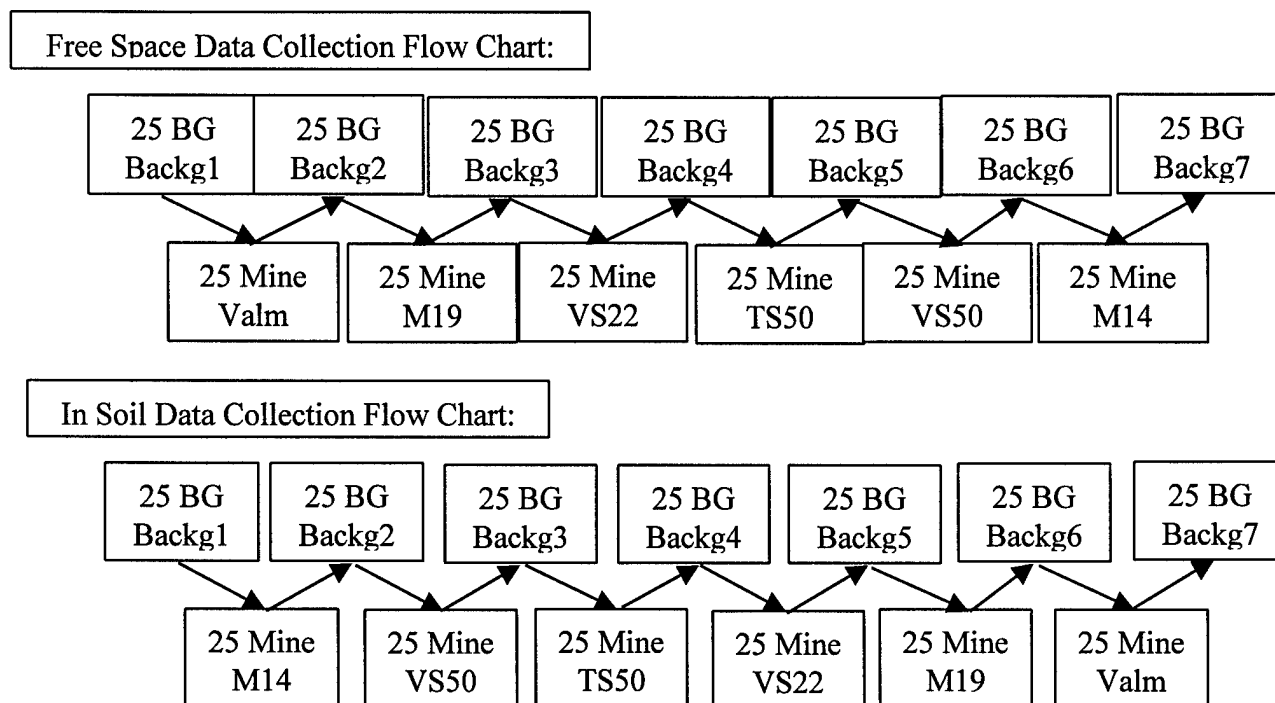
8

Free Space Data Collection Flow Chart:

25 BG Backg1	25 BG Backg2	25 BG Backg3	25 BG Backg4	25 BG Backg5	25 BG Backg6	25 BG Backg7

25 Mine Valm	25 Mine M19	25 Mine VS22	25 Mine TS50	25 Mine VS50	25 Mine M14

In Soil Data Collection Flow Chart:

25 BG Backg1	25 BG Backg2	25 BG Backg3	25 BG Backg4	25 BG Backg5	25 BG Backg6	25 BG Backg7

25 Mine M14	25 Mine VS50	25 Mine TS50	25 Mine VS22	25 Mine M19	25 Mine Valm

Figure 3. Flow chart of data collection in free space and soil. Arrows indicate progression of time. BG = Background measurement/no mine present, Mine = Mine measurement. Second line of text indicates the "scan identifier" associated with the 25-measurement scan. Each scan required approximately 6 minutes to complete.

Prior to repeating the initial experiment, the collected data were analyzed to determine whether the experimental protocol should be modified. An important observation that resulted from this analysis was that the mean, or average, value of the "background" scans varied as a function of time, even in free space. Hereafter, this phenomenon is referred to as "sensor drift". Based on this, the number of samples measured for each scan was reduced to 25. This reduced the amount of time required to complete a scan, and thus reduced the amount of drift that occurred within a scan. This was important because, as will be discussed later, it had previously been assumed that the average of the measurements performed during the course of a scan was a valid estimate of the mine or background signature. This assumption proved to be untrue. The underlying noise process is not stable over time. This issue will be the subject of future analyses.

In addition, after consulting with Geophex engineers, the frequency range over which data was taken was modified to reflect the range used in their initial experiment in which the low-metal content mines were observed to be different than the background. Thus, the second data collection was identical to the first except for the following:

(1) 25 measurements were performed in each scan instead of 50.
(2) The frequency range spanned 90 Hz to 24,000 Hz in 29 logarithmic steps.

E. Sensor/Target Configuration – Final Data Collection

Data was collected under two conditions: a) with the sensor situated in "free-space" (*i.e.* suspended in the air) at a fixed distance from the targets, and b) with the sensor situated at a fixed distance above the ground in which the targets were buried to a depth of 1". Under each condition, data was gathered alternatively with the target absent (background scan), then with the target present (see Figure 3). The frequencies that were used, 19 total, in the second data collection are listed in Table II in the third column. These frequencies were generated in a range from 90 Hz to 24 kHz using a logarithmic spacing. For each target and background scan 25 measurements were made.

Data was gathered first in the free-space condition. For all of the mines except the Valmara, the distance between the top of the mine and the bottom of the sensor head was approximately 10 cm. These distances are listed in Table III. The sensor to mine distance for the Valmara had to be increased since, due to the higher metal content, the response that was exceeded the dynamic range of the sensor at the closer distance.

Next, data was gathered with the mines buried 1 inch from the top of the mine to the surface of the soil. Again, for all of the mines except the Valmara, the distance between the top of the mine and the bottom of the sensor head was approximately 10 cm (see Table III). The data was acquired in the reverse order from that in the free-space condition in order to disturb increasing amounts of soil as the data collection progressed (see Figure 3).

Mine	Target/Mine distance – air (cm)	Target/Mine distance – soil (cm)	File name
Valmara	28	27.5	valm.ppm
M19	11	10.5	m19.ppm
VS2.2	10	9.5	vs22.ppm
TS50	9.5	10	ts50.ppm
VS50	9.5	10	vs50.ppm
M14	10	9.5	m14.ppm

Table III. Experimental parameters used in second experiment.

Following each data collection, data was stored in an ASCII file on the hard disk. The subdirectory of the file indicates the test condition (*e.g.,* 'air' or 'soil'), and the name of the file denotes the scan (scan_identifier). The background conditions are named sequentially as backg1.ppm, backg2.ppm, ... backg7.ppm (see also Figure 3). The names of the files associated with each of the mines are self-explanatory, but are listed in Table III. These files were converted to similarly named files with a .dat extension in which the header has been removed. These .dat files are contained in the appropriate subdirectories on the disk accompanying this report.

For the measurements taken in free-space, the GEM-3 was mounted on a wooden rack with the sensor head approximately 6 feet from the wooden base of the platform (see Figure 4). Both platform and rack were manufactured with no metal parts. The rack assembly allowed placement of targets on a wooden shelf at various distances beneath the sensor head. The sensor head was centered over the mine manually for each measurement. The GEM-3 was linked via cable to the serial port of the laptop PC located approximately 30 feet from the test assembly (see Figure 5).

Figure 4. GEM-3 mount and support system for free-space data collection. Sensor head is located approximately 6 feet from the ground.

Figure 5. The GEM-3 was controlled remotely from a laptop PC. The free-space assembly is behind and to the right of the picnic table (out of view). The soil assembly is behind and to the left of the point from which the picture is taken.

For the measurements taken in soil, the GEM-3 was mounted on a combination of a sawhorse and wooden stand (see Figure 6). The radial location of the wooden stand was adjusted to raise and

lower the sensor head to achieve the desired distance from the ground. The sensor head was centered over the target manually for each measurement.

Figure 6. GEM-3 mount and support system for in-soil data collection.

III. Analysis and Results

Several analyses were performed in order to investigate the statistics of the sensor noise process as well as to extract the signatures of the mines from the measured data. These will be described in the following subsections.

A. General Approach and Notation

This section describes the basic mathematical manipulations that were performed on the raw data measured by the GEM-3. We first describe the normalization required to transform the raw data into data proportional to the induced magnetic field strength. We then define notation that will allow us to refer to specific measurements taken in a particular scan.

In order to extract the signatures of the mines from the measured data, we assume that the best estimate of the signal represented by the set of measurements comprising a scan is the average of those measurements. The notation for this analysis is also defined. We further assume that when the average of the set of measurements comprising the background scans taken before and after a mine scan is subtracted from the average of the mine scan, the result is an estimate of the mine's signature. Below, we specifically define this relationship for this "background corrected" signature.

To evaluate the statistics of the sensor noise process, the standard deviation of the set of measurements comprising a scan is defined. This measure gives an estimate of the variability of

12

the signature. Finally, the set of data acquired during the course of a scan at a particular frequency is defined. This data will be used to estimate the pdfs of the noise process.

(1) As discussed previously, the GEM-3 records in-phase and quadrature data as a function of frequency in ppm. These data are based on a current measurement, as opposed to a magnetic field measurement. Therefore, to convert these measurements to data that are proportional to magnetic field strength, an inverse dependence on the frequency and a 90-degree phase shift must be incorporated (see Section **II.C** for details). Thus, the following calculations were used for each of the measurements taken:

$$I(f) = -\frac{Q_s(f)}{f}$$
$$Q(f) = \frac{I_s(f)}{f} \quad .$$

Here, f is frequency, $I_s(f)$ and $Q_s(f)$ are the in-phase and quadrature samples reported by the GEM-3, and $I(f)$ and $Q(f)$ are the calculated in-phase and quadrature samples that are proportional to the magnetic field.

(2) The in-phase and quadrature responses obtained on the i^{th} measurement as a function of frequency are denoted either $I^i_{scan\_identifier}(f)$ or $Q^i_{scan\_identifier}(f)$. Here, i ranges from 1 to 25 (or 50), and *scan_identifier* refers to the scan, as well as the file name that was used to store the data (*e.g.,* backg3, M19, etc.). For example, $I^9_{M14}(f)$ is the in-phase data measured at each frequency on the 9^{th} (of 25) measurement during the M14 scan.

(3) The mean in-phase and quadrature signatures for a particular scan are calculated as the average of all of the 25 (or 50) measurements via

$$\bar{I}_{scan\_identifier}(f) = \frac{1}{25}\sum_{i=1}^{25} I^i_{scan\_identifier}(f) \text{ and}$$

$$\bar{Q}_{scan\_identifier}(f) = \frac{1}{25}\sum_{i=1}^{25} Q^i_{scan\_identifier}(f).$$

The mean signatures are indicated with a bar and do not have a measurement index superscript. Thus, $\bar{I}_{M14}(f)$ is the mean of the in-phase measurements taken during the M14 scan as a function of frequency.

(4) A "background corrected" average signature is derived for each mine by subtracting an estimate of the background signature from the mean of the scan for each mine. The background estimate is calculated as the average of the background measurements taken immediately prior to and after the mine scan. For example, the background estimate for the M14 mine in free space would consist of the average of the measurements taken in the 6^{th} and 7^{th} background scans. The background corrected signature is denoted by a '*bc*' in the subscript, and is given by

$$\bar{I}_{scan\_identifier,bc}(f) = \bar{I}_{scan\_identifier}(f) - \frac{\bar{I}_{previous\_bg\_identifier}(f) + \bar{I}_{next\_bg\_identifier}(f)}{2} \text{ and}$$

$$\bar{Q}_{scan\_identifier,bc}(f) = \bar{Q}_{scan\_identifier}(f) - \frac{\bar{Q}_{previous\_bg\_identifier}(f) + \bar{Q}_{next\_bg\_identifier}(f)}{2}.$$

13

We also define the average background signature *for each mine* as the second term in the above equations, or

$$\bar{I}_{scan\_identifier}^{BACKGROUND}(f) = \frac{\bar{I}_{previous\_bg\_identifier}(f) + \bar{I}_{next\_bg\_identifier}(f)}{2}$$

$$\overline{Q}_{scan\_identifier}^{BACKGROUND}(f) = \frac{\overline{Q}_{previous\_bg\_identifier}(f) + \overline{Q}_{next\_bg\_identifier}(f)}{2}$$

The superscript '*BACKGROUND*' indicates that this is an average background signature associated with the mine denoted by the scan identifier. Rewriting the first two equations using the above notation,

$$\bar{I}_{scan\_identifier,bc}(f) = \bar{I}_{scan\_identifier}(f) - \bar{I}_{scan\_identifier}^{BACKGROUND}(f)$$

$$\overline{Q}_{scan\_identifier,bc}(f) = \overline{Q}_{scan\_identifier}(f) - \overline{Q}_{scan\_identifier}^{BACKGROUND}(f)$$

For the M14 in free space,

$$\bar{I}_{M14,bc}(f) = \bar{I}_{M14}(f) - \frac{\bar{I}_{backg6}(f) + \bar{I}_{backg7}(f)}{2} = \bar{I}_{M14}(f) - \bar{I}_{M14}^{BACKGROUND}(f) \text{ and}$$

$$\overline{Q}_{M14,bc}(f) = \overline{Q}_{M14}(f) - \frac{\overline{Q}_{backg6}(f) + \overline{Q}_{backg7}(f)}{2} = \overline{Q}_{M14}(f) - \overline{Q}_{M14}^{BACKGROUND}(f).$$

(5)   The standard deviation of the in-phase and quadrature signatures is calculated for each scan as

$$\sigma_{scan\_identifier}^{I}(f) = \sqrt{\frac{1}{24}\sum_{i=1}^{25}(I_{scan\_identifier}^{i}(f) - \bar{I}_{scan\_identifier}(f))^2} \text{ and}$$

$$\sigma_{scan\_identifier}^{Q}(f) = \sqrt{\frac{1}{24}\sum_{i=1}^{25}(Q_{scan\_identifier}^{i}(f) - \overline{Q}_{scan\_identifier}(f))^2}.$$

For example, $\sigma_{M14}^{I}(f)$ is the standard deviation as a function of frequency of the in-phase measurements taken during the M14 scan.

(6)   The set of in-phase or quadrature samples taken during a particular scan at a given frequency, f_0, are denoted $I_{scan\_identifier}^{\{i\}}(f_0)$ and $Q_{scan\_identifier}^{\{i\}}(f_0)$ respectively, and are given by the sets

$$I_{scan\_identifier}^{\{i\}}(f_0) = \{I_{scan\_identifier}^{1}(f_0), I_{scan\_identifier}^{2}(f_0), \ldots I_{scan\_identifier}^{25}(f_0)\} \text{ and}$$

$$Q_{scan\_identifier}^{\{i\}}(f_0) = \{Q_{scan\_identifier}^{1}(f_0), Q_{scan\_identifier}^{2}(f_0), \ldots Q_{scan\_identifier}^{25}(f_0)\}$$

Thus, $I_{M14}^{\{i\}}(900) = \{I_{M14}^{1}(900), I_{M14}^{2}(900), \ldots I_{M14}^{25}(900)\}$ is the set of in-phase measurements taken during the M14 scan at 900 Hz.

B. Characterization of the Noise Process

The goal of this analysis is to characterize the statistics of the noise. These statistics will be used to formulate the detection algorithm prescribed by signal detection theory. The probability density function (pdf) that describes the noise process (*e.g.*, Gaussian) was evaluated by forming histograms of the data taken during a scan at each frequency (see **III.A.6** above). We generated these histograms in both soil and free-space. In general, 50 samples of a noise process are adequate to obtain estimates of its mean and variance. However, 50 is not a large enough sample to accurately estimate the functional form of a pdf, or histogram. We provide an example of this in the Results section). Therefore, as part of a future experiment, a 1000 measurement scan of a "clean" background will be performed in order to obtain a better estimate of the functional form of the pdf, as well as to address the sensor drift issue in more detail.

In order to determine whether or not the statistics of the noise process are stable, we examined the mean and variance of the statistical distributions as a function of the following parameters:

    (1) the absolute amplitude level of the received signal
    (2) the amount of time the sensor has been operating, corresponding to something akin to "drift" in the mean response of the sensor, and
    (3) the frequency at which a measurement is made.

In the initial experimental design, 50 measurements were to be made in each scan. This data was to be used to 1) analyze the statistics of the noise process, and 2) to determine the detectability of low metal mines in both free space and when buried in soil. Although our original goal was to characterize the pdfs of the noise process by analyzing all of the background data together (a total of 350 measurements in both the free space and soil conditions), the sensor drift issue made this infeasible. Figure 7 illustrates the sensor drift problem using histograms of the 2,370 Hz data measured in each of the free space background scans. The histogram for each scan, backg1 through backg7, follows the color scheme listed in the legend. In general, the mean of each scan is increasing with time (scan number), however the mean of backg5 is less than the means of backg3 and backg4. This indicates that the mean, as well as the standard deviation, of the sensor response is not stationary in time. Clearly, combining data across all scans cannot generate a valid estimate of the pdf of the noise process.

15

Figure 7. Color-coded histograms of the data obtained at 2,370 Hz during the free space background measurements. The change in both the mean and standard deviation of each data set indicates that the sensor response "drifts", or is not stable over time.

Although 50 data points are sufficient to provide relatively accurate estimates of the mean and variance of a process, 50 points are not sufficient to differentiate between a set of candidate pdfs. To illustrate this point, we used a random number generator to produce 50 samples from a Gaussian distribution and 50 samples from a uniform distribution. Histograms of these data are provided in Figure 8. Clearly, it is difficult to definitively say that the histogram of one set of data is indicative of a Gaussian distribution and the other a uniform distribution.
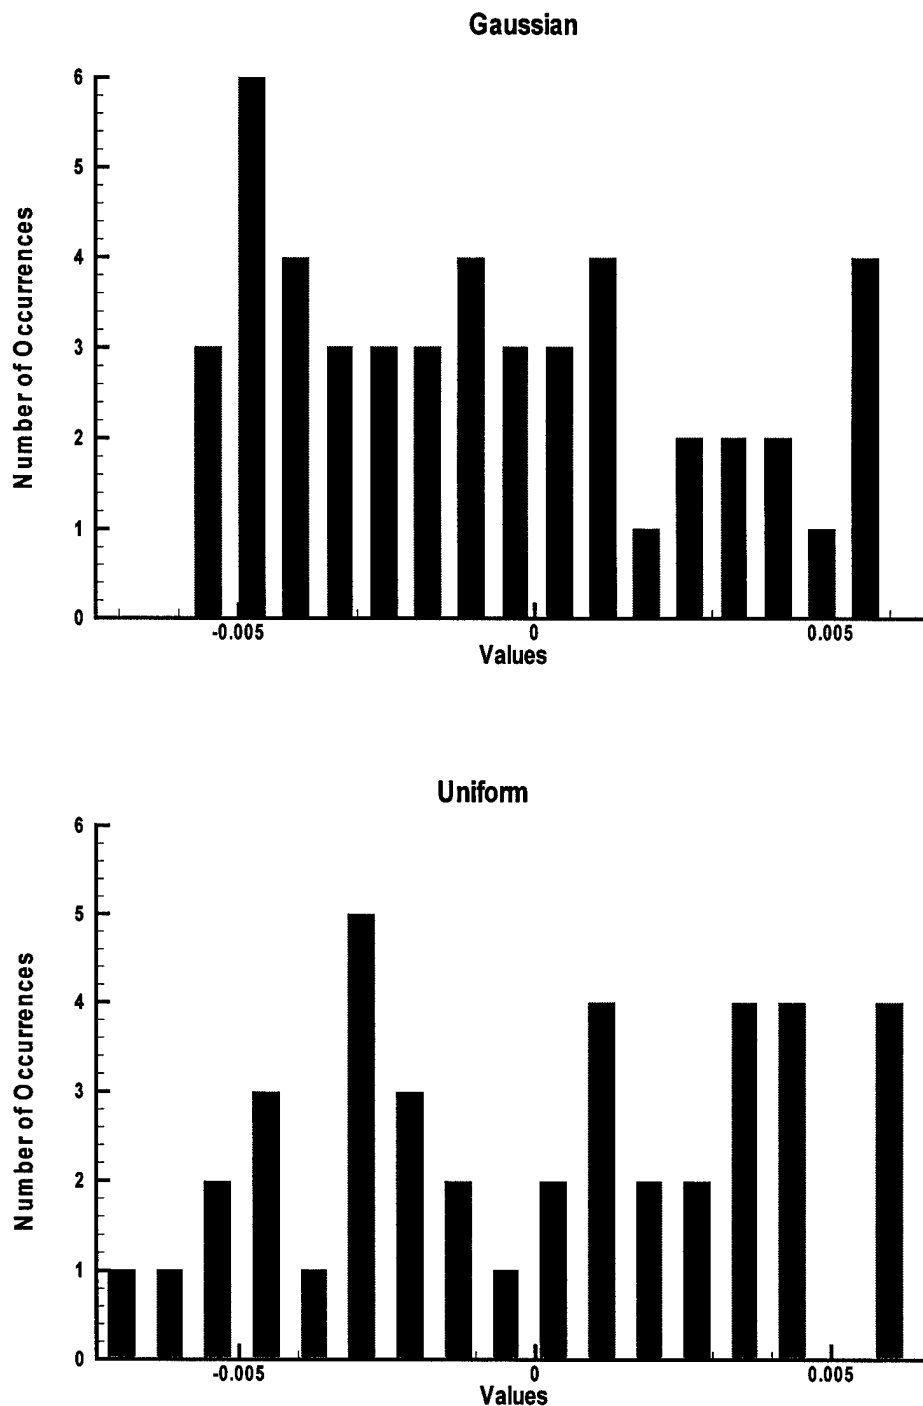
Figure 8 – Histograms of 50 samples of a Normal/Gaussian and Uniform density function.

In order to investigate the sensor drift issue further, Figures 9 and 10 plot the average response as a function of frequency for each of the background measurements (backg1 – backg7) for free-space and soil, respectively. These are based on the experimental data in which 25 measurements were taken. Since the data below 300 Hz was extremely noisy, it is not plotted. The noisiest (highest variance) data occurs at integers of 60 Hz. Figures 11 and 12 plot the standard deviation of the 25 measurements as a function of frequency for each of the background measurements

taken in free-space and soil, respectively. Both Figure 8 and Figure 9 indicate that the mean response of the sensor changes substantially as a function of time – indicating that the mean is not stationary. Comparing Figure 9 with Figure 8 indicates that variability in the density or moisture level of the soil (associated with digging) may also change the "background" signature recorded by the sensor. However, based on this analysis, it is difficult to separate out the differences that result from changes in the soil density from differences that result from slight changes in the sensor to soil distance. This effect should be considered for advanced algorithm development. Figures 11 and 12 indicate that the level of the variability (standard deviation) is a function of frequency, and also appears to be a function of time. In general, more variability is observed at high and low frequencies than in the middle of the frequency range. Clearly, this will have to be considered for algorithm development.
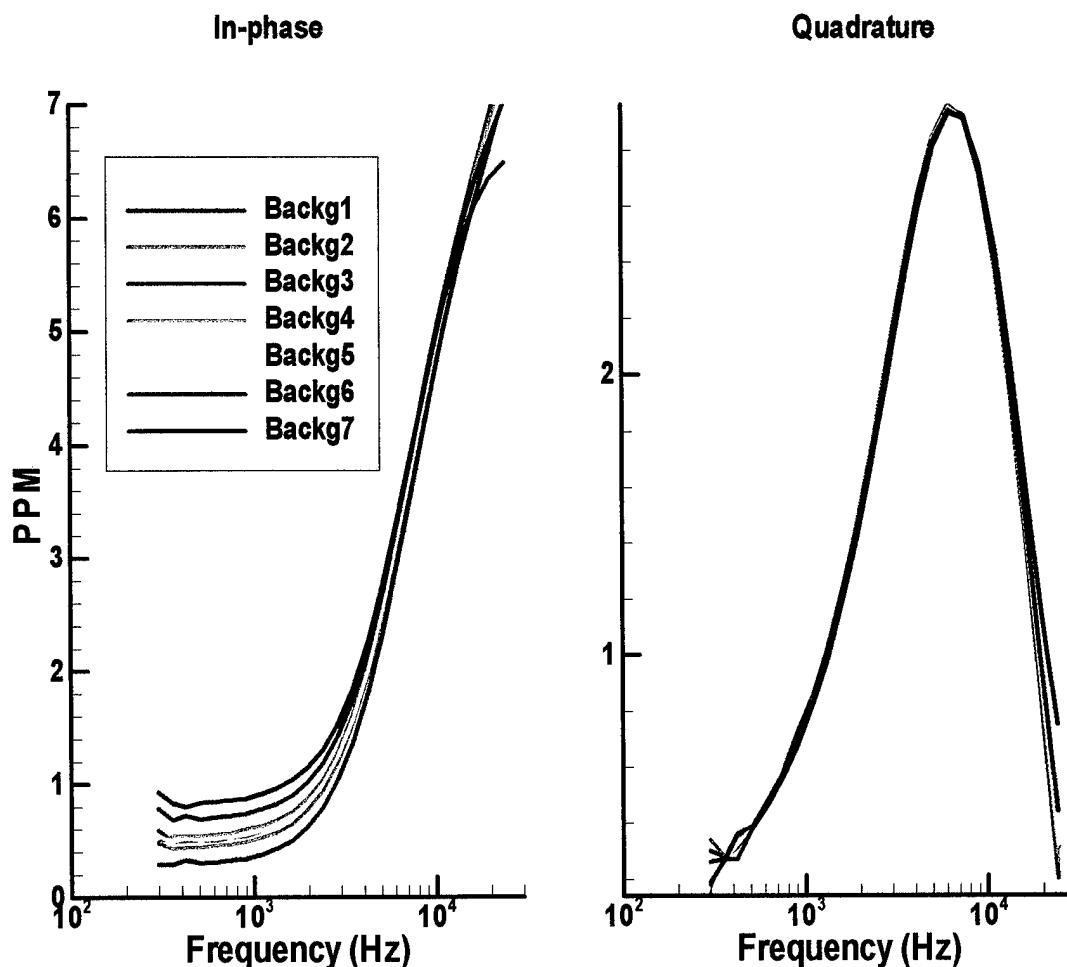
Figure 9 – Mean in-phase and quadrature background variation for the seven measured samples in free-space. Color order (beginning with backg1 and ending with backg7) = red, green, blue, cyan, yellow, magenta, black.

18

Figure 10 – Mean in-phase and quadrature background variation for the seven measured samples in soil. Color order (beginning with backg1 and ending with backg7) = red, green, blue, cyan, yellow, magenta, black. Note: bacgk7.ppm (black line) measured at a different height.
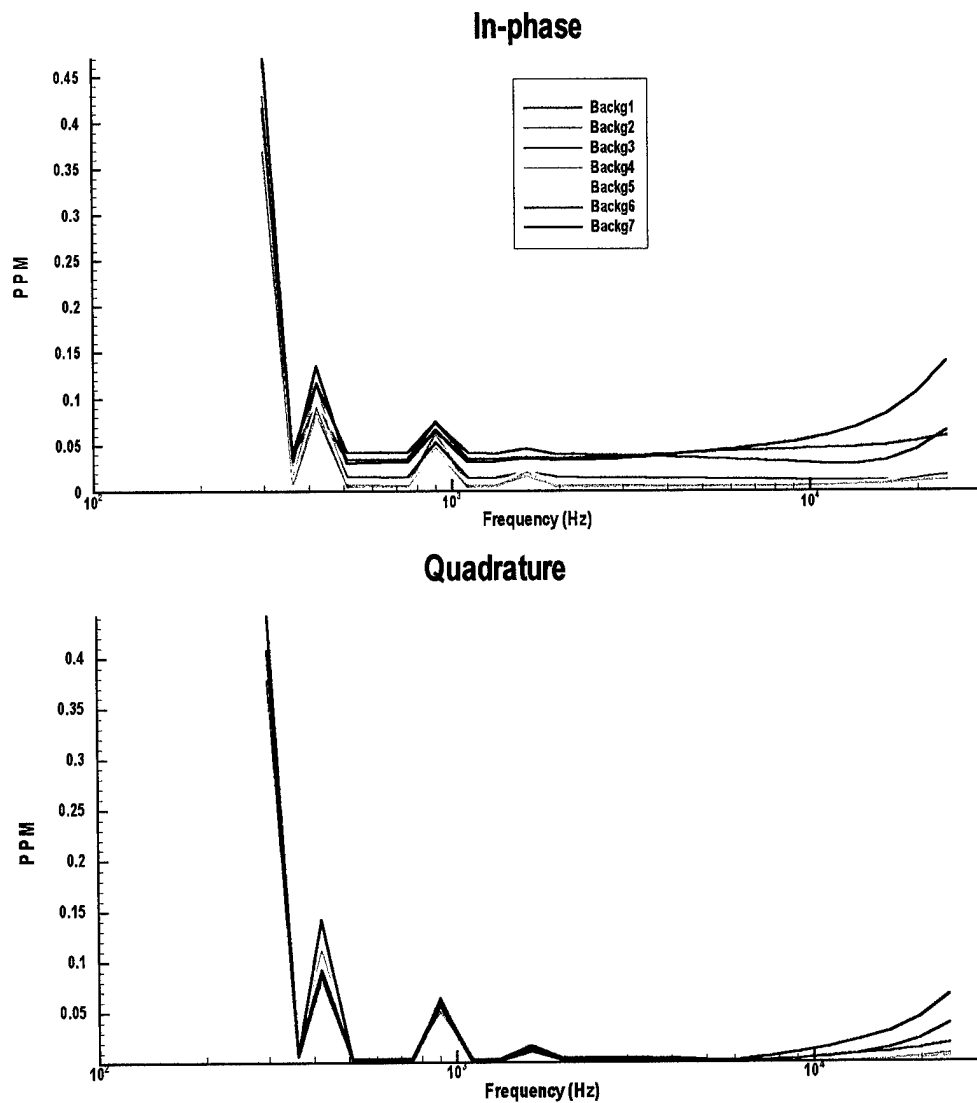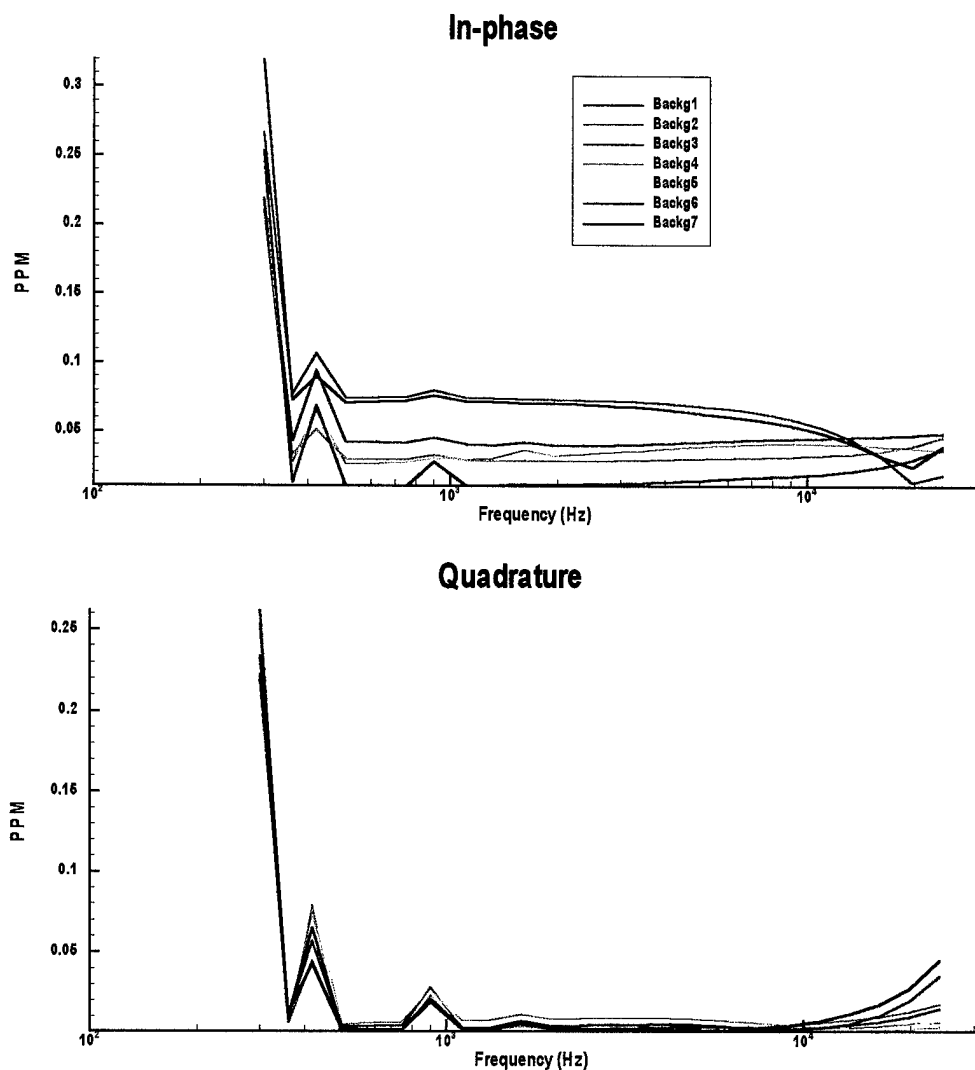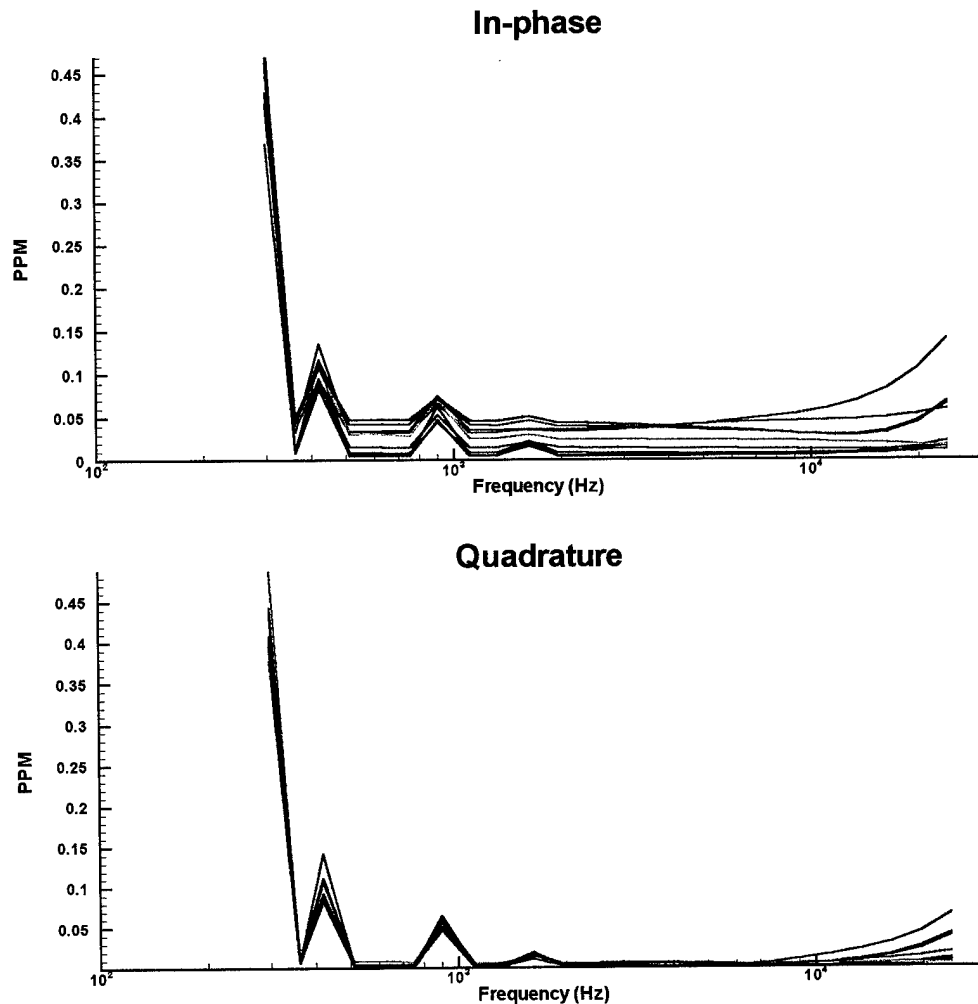
Figure 11 – Standard deviation of the data as a function of frequency in free-space from each of the seven measurements. Color order (beginning with backg1 and ending with backg7) = red, green, blue, cyan, yellow, magenta, black.

Figure 12 – Standard deviation of the data as a function of frequency in soil from each of the seven measurements. Color order (beginning with backg1 and ending with backg7) = red, green, blue, cyan, yellow, magenta, black.

In-phase

PPM

0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0

10^2 10^3 10^4

Frequency (Hz)

Quadrature

PPM

0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05

10^2 10^3 10^4

Frequency (Hz)

Figure 13 – Standard deviation of the data as a function of frequency in free space from each of the seven background measurements (shown in blue), from the M14 (shown in red), and from the Valmara (shown in green).

In order to evaluate the effects of absolute amplitude of the received signal on the noise statistics, the standard deviations observed in the background scans were compared to those associated with the M14 and the Valmara. Figure 13 plots the standard deviation of the free-space background scans (shown in blue) along with the standard deviation of the M14 scan (shown in red) and the Valmara scan (shown in green). Clearly, there is no definitive relationship between signal strength and noise strength, indicating that the noise process is stationary with respect to signal amplitude.

Figures 14 and 15 plot each of the measured background signatures obtained in free-space and soil, respectively, and also serve to indicate the extent of the drift. Additionally, it can be observed from Figure 15 that as more and more soil is disturbed, the quadrature component of the measured response increases. No clear pattern can be observed in the in-phase component. Recall, the order of burial was M14, VS50, TS50, VS2.2, M19, Valmara and background measurements were made following each emplacement. All soil was returned and tamped down once the mine was removed, but the measurement was taken over the disturbed soil.
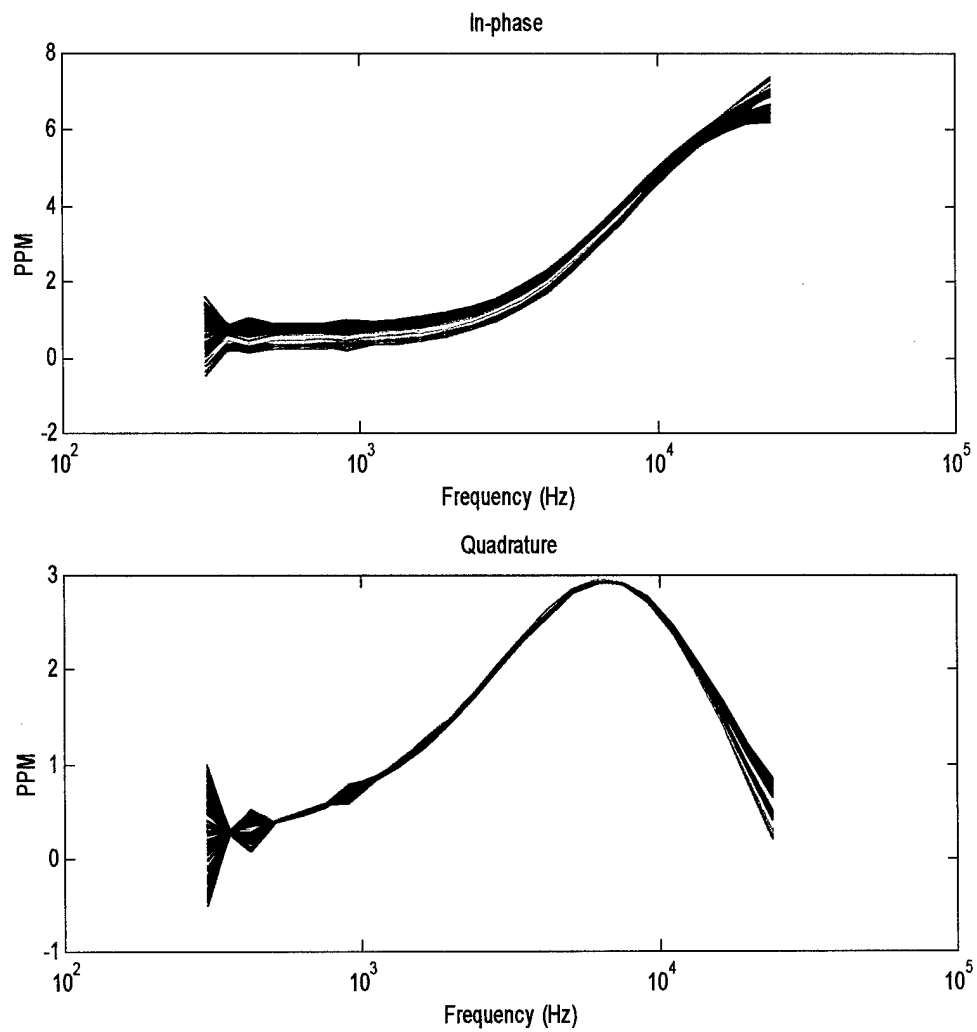
Figure 14 – All in-phase and quadrature background signatures for the seven measured samples in free-space. Color order (beginning with backg1 and ending with backg7) = blue, green, red, cyan, yellow, magenta, black.
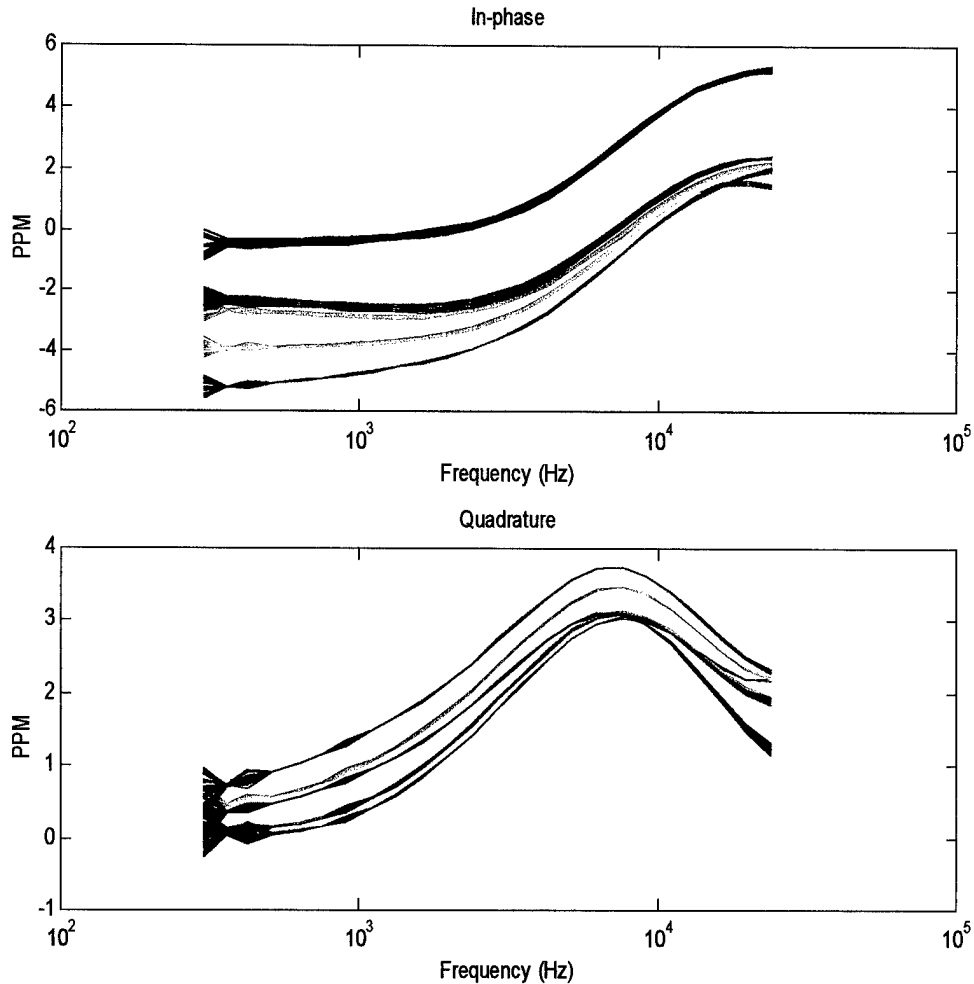
In-phase

Quadrature

Figure 15 – All in-phase and quadrature background signatures for the seven measured samples in free-soil. Color order (beginning with backg1 and ending with backg7) = blue, green, red, cyan, yellow, magenta, black. Note: bacgk7.ppm (black line) measured at a different height.

C. Calculation of the Average Background Corrected Mine Signatures

The average of each scan was calculated as described in **III.A.3**. To calculate the response of the sensor to the mines without including the background response, the background response must be subtracted from the response obtained with mine plus background. In free space, this was performed by subtracting an estimate of the background signature measured before and after each mine measurement (see **III.A.4**). In soil, the same approach was taken, with the exception of the Valmara 69. When the Valmara scan was performed in soil, the distance between the sensor and the ground was larger than the distance used for the other mines. Consequently, the background signature measured following the Valmara scan (backg7) was taken with the sensor at the same (greater) distance from the ground. Therefore, for the Valmara scan in soil, the background-corrected signature was calculated as

$$\bar{I}_{valm,bc}(f) = \bar{I}_{valm}(f) - \bar{I}_{backg7}(f)$$
$$\bar{Q}_{valm,bc}(f) = \bar{Q}_{valm}(f) - \bar{Q}_{backg7}(f)$$

As will be discussed further in the Results section, the background changed substantially after the mines were emplaced. Therefore, in addition to calculating the background corrected signals as described in **III.A.4**, mine signatures were also calculated by subtracting the initial background measurement, backg1. This scan was taken prior to disturbing the soil. The drawback to this approach is that it does not incorporate any effects of sensor drift.

Figure 16 shows the average background corrected signatures for the six mines in free-space. Each of the mines appears to have a different characteristic signature as a function of frequency. Again, note that no clutter signatures were measured for comparison. Figure 16 also illustrates the need for accurate background correction. As discussed in the previous section, the mean sensor response that is measured when no objects are present "drifts" over time. The level of this drift can be as high as 1 ppm over the course of a 2-3 hour period (see Figures 7, 9, and 14). The ppm levels of the background corrected signatures for some of the low metallic-content mines are smaller than this level of drift. Therefore, correcting a measured signal with an "old" background measurement could substantially change the form of the background-corrected signal. Clearly, this statistical variability must be built into detection algorithms.

Figure 17 shows the background corrected signatures for the six mines buried 1" deep in soil. Again, each of the mines appears to have a different characteristic signature as a function of frequency. However, for the low-metal mines, these are not the same signatures as those measured in free-space, although the signatures for the Valmara and the VS50 match quite well.

Figure 18 plots the results shown in Figures 16 and 17 together. Clearly, the signatures measured in the ground are different from those measured in the air. In addition, the signatures measured in the ground have a higher magnitude than those measured in free-space. This may be a result of the GEM-3 measuring the dielectric discontinuity between the mine and the ground. Dr. Lawrence Carin, at Duke University, is attempting to address this hypothesis via a modeling study. Figure 19 plots results similar to those of Figure 18, except that the soil signatures are corrected by the background measured initially - *i.e.* to undisturbed soil.
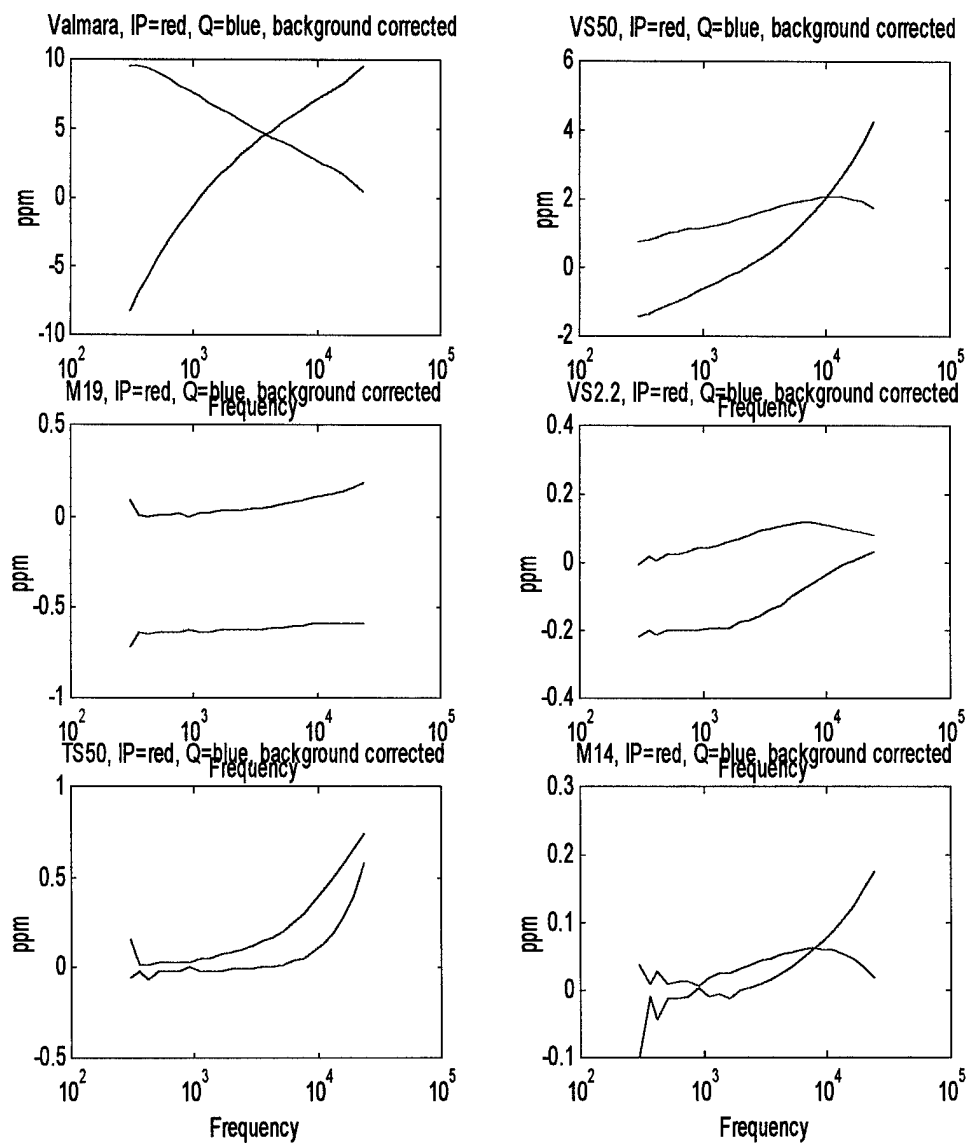
Figure 16 – Background corrected signatures as a function of frequency in free-space.
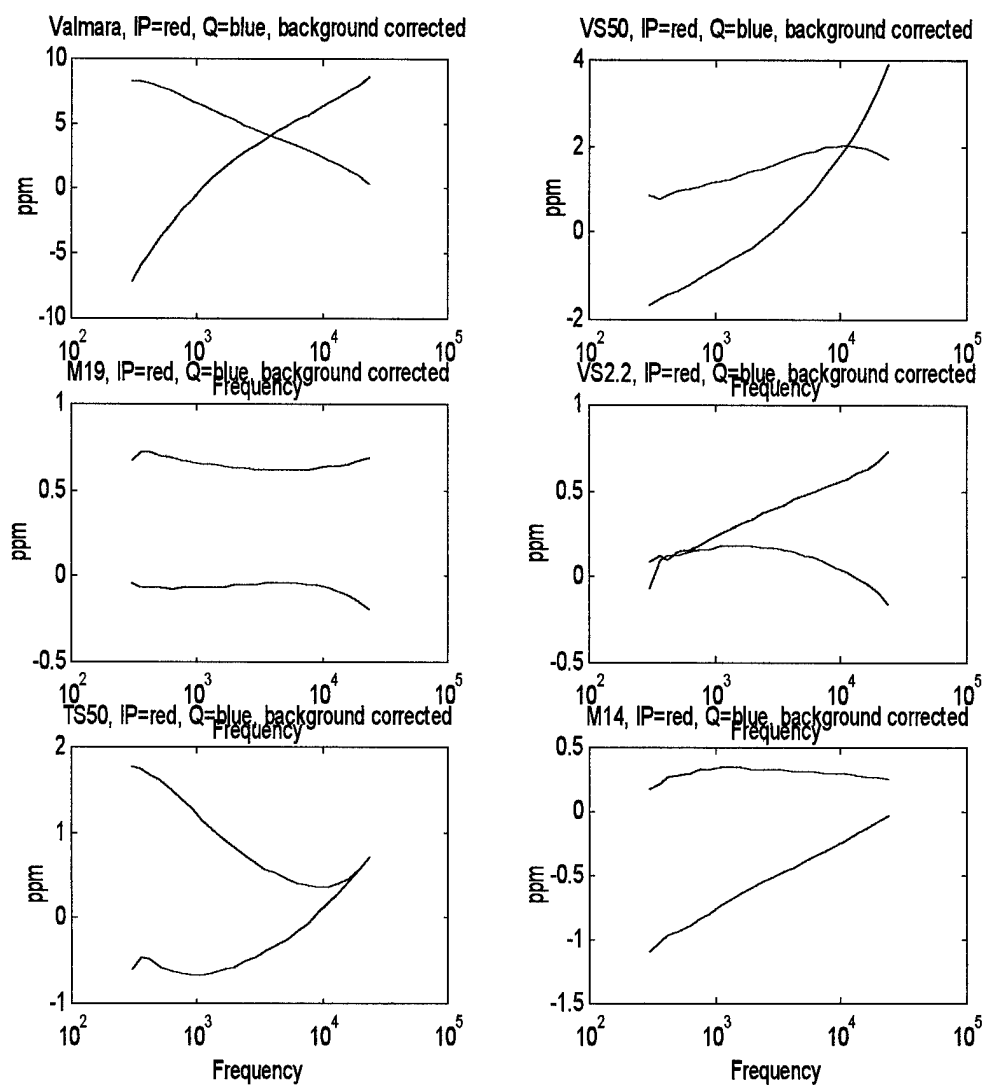
Figure 17 – Background corrected signatures as a function of frequency in soil.
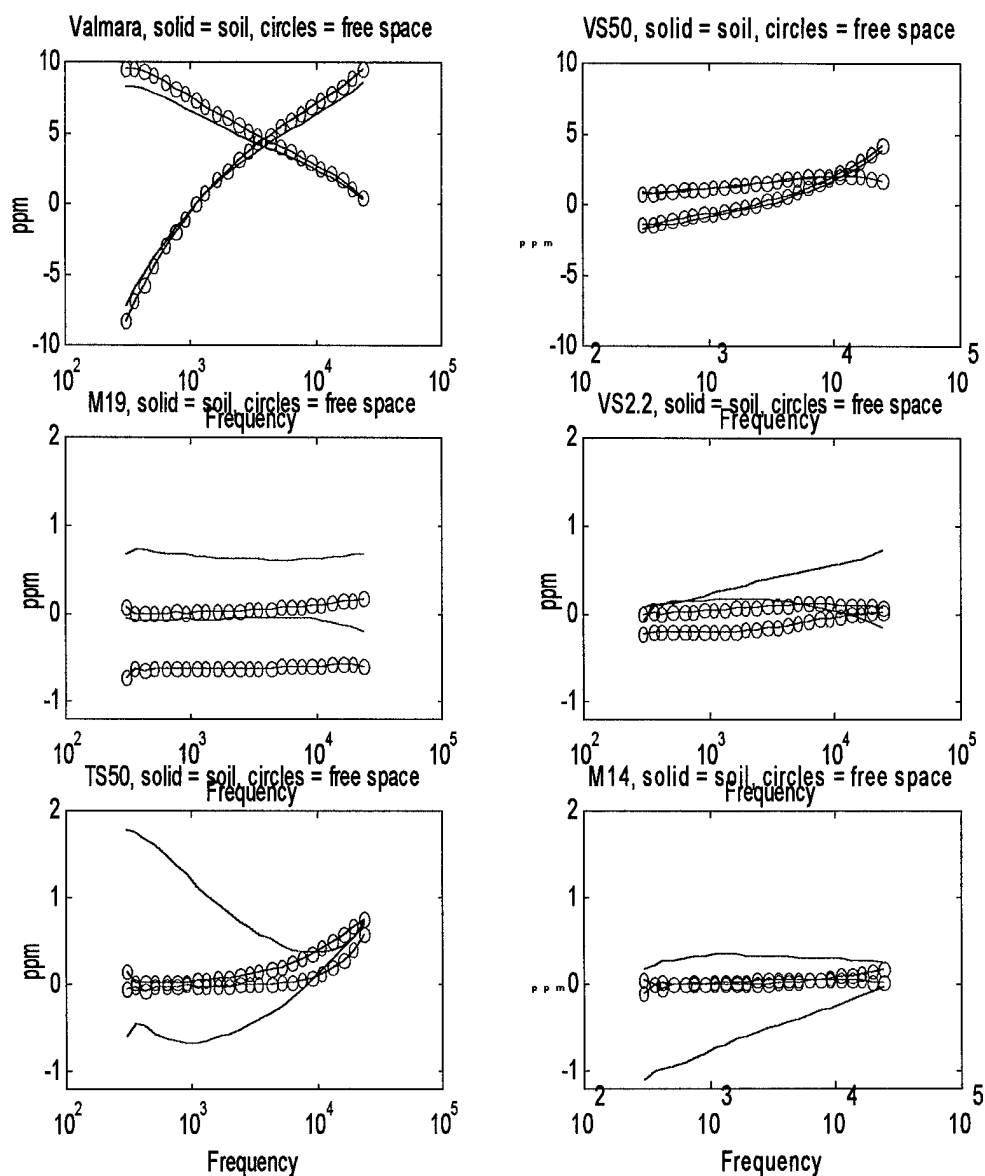
Figure 18 –Comparison of background corrected signatures in air and soil. Red = in-phase component, blue = quadrature component, solid lines measured in soil, solid lines with circles measured in free-space.
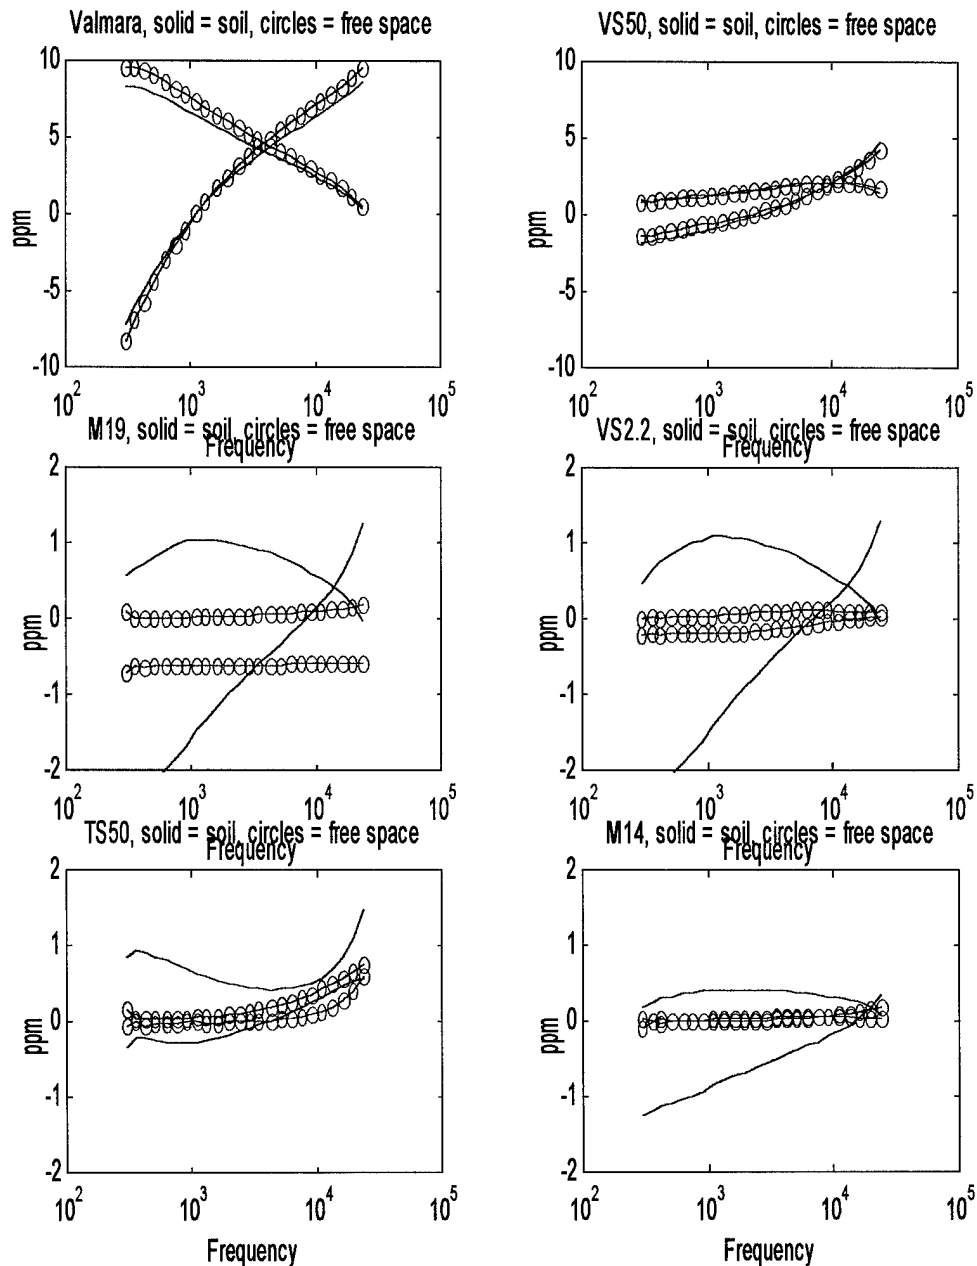
Figure 19 –Comparison of background corrected signatures in air and soil. Soil signatures corrected by initial background measurement. Red = in-phase component, blue = quadrature component, solid lines measured in soil, solid lines with circles measured in free-space.

D. Detectability analysis

The goal of these analyses was to determine whether the sensor responses obtained from the mines were different from the sensor responses obtained from the background (either soil or free-space). Detection performance was measured using Receiver Operating Characteristic (ROC) analysis. In order to employ ROC analysis, the following are required. First, a set of data is needed in which repeated measurements were taken under both hypotheses: mine present and

mine absent. This data is available in the mine scans and background scans, respectively. For example, the detectability of the M14 in free space could be analyzed using the data collected in the M14 scan and data collected in the backg6 and backg7 scans. Second, an algorithm for processing the data must be specified. In this analysis, we have employed two different algorithms, each of which we consider to be a "baseline" algorithm. These algorithms are categorized as "baseline" because they are fairly routine approaches to processing such data.

In order to perform the ROC analysis, we have assumed that the algorithm has at its disposal a set of measured in-phase and quadrature data on which to make a decision. The detector, or algorithm, will not know *a priori* whether or not a mine was present when the data was measured, and must therefore make this decision based on the measured data. We further assume that the decision is made after performing a background correction on the measured data. This would result in a target hypothesis, H_1, in which the background-corrected signal consists of the mine signature plus noise and the null hypothesis, H_0, in which the background-corrected signal consists of noise alone. Therefore, to perform this analysis on the data collected in each set of mine/background scans, a "background" must be subtracted from each measurement taken in the scan. In this analysis, we have chosen to subtract the average background signature associated with each mine scan as defined in **III.A.4**.

For the target hypothesis (H_1), the data available to the detection algorithms consists of the 25 measurements taken in a particular mine scan corrected by the appropriate average background:

$$I^i_{scan\_identifier,bc,H1}(f) = I^i_{scan\_identifier}(f) - \bar{I}^{BACKGROUND}_{scan\_identifier}, i = 1...25,$$

$$Q^i_{scan\_identifier,bc,H1}(f) = Q^i_{scan\_identifier}(f) - \bar{Q}^{BACKGROUND}_{scan\_identifier}, i = 1...25.$$

Thus, the detection algorithms have twenty-five signals available that were obtained when H_1 is true. Note these signals are different than the previous definition of a background corrected signature since each individual measurement has the background removed separately. In the previous definition, the mean signature for the mine was corrected. The notation differs in that, for the above signals, there is no bar associated with the signal and an index superscript, *i*, is used. The affiliation with the H_1 data set is also noted in the subscript. For the null hypothesis (H_0), the data available to the detector consists of the 25 background measurements taken prior to the mine scan along with the 25 measurements taken immediately after the mine scan, each corrected by the appropriate average background:

$$I^i_{scan\_identifier,bc,H0}(f) = I^i_{previous\_bg\_identifier}(f) - \bar{I}^{BACKGROUND}_{scan\_identifier}, i = 1...25,$$

$$Q^i_{scan\_identifier,bc,H0}(f) = Q^i_{previous\_bg\_identifier}(f) - \bar{Q}^{BACKGROUND}_{scan\_identifier}, i = 1...25$$

$$I^{i+25}_{scan\_identifier,bc,H0}(f) = I^i_{next\_bg\_identifier}(f) - \bar{I}^{BACKGROUND}_{scan\_identifier}, i = 1...25$$

$$Q^{i+25}_{scan\_identifier,bc,H0}(f) = Q^i_{next\_bg\_identifier}(f) - \bar{Q}^{BACKGROUND}_{scan\_identifier}, i = 1...25$$

A total of 50 signals that were obtained when H_0 is true are available for processing. The superscripts on the left-hand side of the last two equations use the notation *i+25*. This allows the set of background corrected data to be indexed over the range of 26 to 50 as the measurement index, *i*, ranges from 1 to 25.

The performance of two algorithms was analyzed. The first algorithm that was implemented was an energy detection algorithm. Energy in a real signal, *G(f)*, is calculated as

$$E = \int_{-\infty}^{\infty} G(f)^2 \, df \,.$$

Energy in a complex signal, $C(f) = R(f) + jI(f)$, is calculated as

$$E = \int_{-\infty}^{\infty} |C(f)|^2 \, df = \int_{-\infty}^{\infty} [R(f)^2 + I(f)^2] \, df$$

where $R(f)$ is the real part of $C(f)$, $I(f)$ is the imaginary part of $C(f)$, and j is the square root of -1. In the case where the data consists of samples, the total energy is calculated by squaring each of the sample values and summing them. In our data, there are 29 frequency samples, and the in-phase and quadrature correspond to the real and imaginary parts of the signal, respectively. Thus, the magnitude squared of the k^{th} frequency component, $[M^{i}_{scan\_identifier,bc}(f_k)]^2$, is equal to the in-phase component squared plus the quadrature component squared. The data for the target hypothesis (H_1) was the energy contained in the 25 background-corrected signatures defined above, e.g., for the M14

$$E_{M14\_H1}(i) = \sum_{k=1}^{29} ([I^{i}_{M14,bc,H1}(f_k)]^2 + [Q^{i}_{M14,bc,H1}(f_k)]^2) = \sum_{k=1}^{29} [M^{i}_{M14,bc}(f_k)]^2, i = 1 \dots 25$$

Note the H_1 notation in the subscript to indicate that this set of data is associated with the target hypothesis. The data for the null hypothesis (H_0) was the energy contained in the 50 background-corrected background signatures that were measured immediately prior to and following collection of the target data, e.g.,

$$E_{M14\_H0}(i) = \sum_{k=1}^{29} ([I^{i}_{M14,bc,H0}(f_k)]^2 + [Q^{i}_{M14,bc,H0}(f_k)]^2) = \sum_{k=1}^{29} [M^{i}_{M14,bc}(f_k)]^2, i = 1 \dots 50$$

The second detection algorithm that was implemented was a matched filter. This approach is normally implemented by calculating a matched filter signature either from theoretical considerations or from a training set of measurements, and then implementing the filter on a second set of measurements. However, time constraints did not allow such a design to be implemented in this preliminary experiment. Therefore, for both soil and free space scans, odd-numbered measurements within a mine scan were used to calculate the matched filter. Performance was evaluated on the even-numbered mine measurements and all of the associated background measurements. In addition, the matched filter derived from a free space mine scan was applied to data obtained in soil.

As described above, the matched filter was calculated by taking the average of the odd measurements of the background-corrected target signatures,

$$I_{scan\_identifier,bc}^{matched\_filter}(f) = \frac{1}{13}\sum_{i=1,3,...25} I_{scan\_identifier,bc,H1}^{i}(f),$$

$$Q_{scan\_identifier,bc}^{matched\_filter}(f) = \frac{1}{13}\sum_{i=1,3,...25} Q_{scan\_identifier,bc,H1}^{i}(f)$$

The superscript '*matched filter*' indicates that these are the matched filter signals. A correlation operation on two signals is performed by multiplying the individual frequency components point by point and then summing. The output of the matched filter, m, for the target hypothesis (H_1) is the correlation between this matched filter and the even scans of the background-corrected signatures, *i.e.*,

$$m_{scan\_identifier,H1}(\frac{i}{2}) = \sum_{k=1}^{29}[(I_{scan\_identifier,bc,H1}^{i}(f_k)xI_{scan\_identifier,bc}^{matched\_filter}(f_k) +$$

$$Q_{scan\_identifier,bc,H1}^{i}(f_k)xQ_{scan\_identifier,bc}^{matched\_filter}(f_k), i = 2,4...24$$

The notation $\frac{i}{2}$ is required to correctly index the output of the matched filter from 1 to 12, since the index, i, proceeds through the even integers between 2 and 24. Note that this formulation assumes that the in-phase and quadrature components of the received signal are independent. If the sensor noise is 1) additive, 2) Gaussian, and 3) uncorrelated across time, the assumption of independence is reasonable. The data for the null hypothesis (H_0) was the correlation between this matched filter and the 50 background-corrected background signatures that were measured immediately prior to and following the collection of the target data, *i.e.*,

$$m_{scan\_identifier,H0}(i) = \sum_{k=1}^{29}[(I_{scan\_identifier,bc,H0}^{i}(f_k)xI_{scan\_identifier,bc}^{matched\_filter}(f_k) +$$

$$Q_{scan\_identifier,bc,H0}^{i}(f_k)xQ_{scan\_identifier,bc}^{matched\_filter}, i = 1...50$$

This analysis was performed in both soil and in free-space.

If the GEM-3 sensor responds solely to the metal contained in the mine (*i.e.*, the presence or absence of the soil does not affect the measurement), then identical signatures should be measured in both the free space and soil conditions. In this case, a matched filter derived from the free space measurement could be used to detect the presence of a particular mine in soil. Therefore, in addition to the matched filter analyses described above, the matched filters that were calculated for the free space condition were also applied to the buried targets. This analysis allowed us to determine whether the signatures measured in free space were useful for detecting the same targets in soil.

Figure 20 shows the results of the ROC analysis for the six mines in free space. Figure 21 shows similar results for the six mines in soil. The red curves illustrate the performance of the matched filter detector; the blue lines illustrate the performance of the energy detector. In the where the blue line is not visible cases (*e.g.* the Valmara ROC), the blue line is directly under the red line in the upper left-hand corner. In free space, the energy detector achieves perfect performance on four of the six mines. It is slightly less than perfect for the VS2.2, and its performance falls below the chance diagonal for the M14. This anomaly results because the energy in the background-corrected signature is lower for the M14 than for the background. The matched filter

detector achieves perfect performance on five of the six mines, and substantially outperforms the energy detector on the M14.

In soil, the matched filter detector performance is perfect on all mines. The energy detector performance is perfect on all but the VS2.2, where it always chooses the wrong hypothesis since again, the energy in the background corrected signature is lower when the VS2.2 is present. This, and the above somewhat anomalous result, may be an artifact of a poor background correction, which results from drift in the background. When the mean background corrected signatures from free space were used as a matched filter in the soil, detection performance was still perfect.
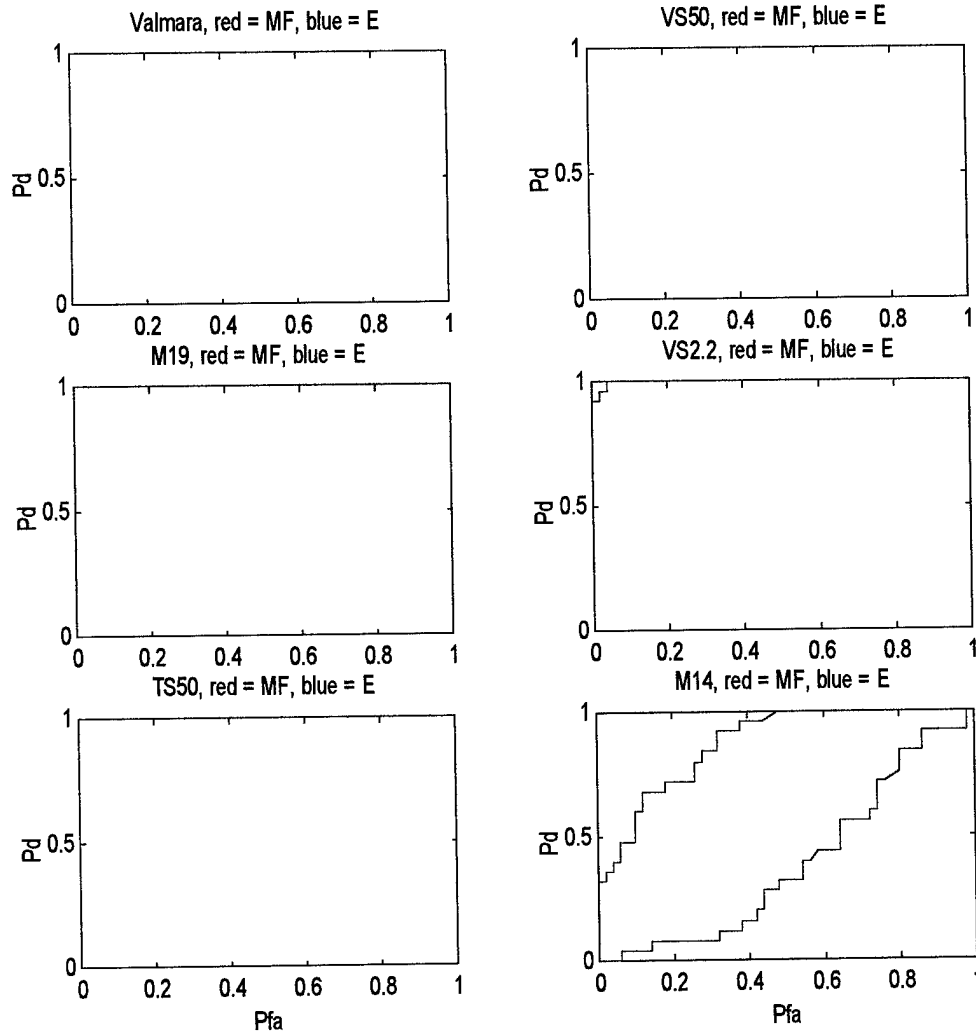

Figure 20 – ROCs for the six mines in free-space. Red lines indicate matched filter detector performance, blue lines indicate energy detector performance.

Figure 21 – ROCs for the six mines in soil. Red lines indicate matched filter detector performance, blue lines indicate energy detector performance.

E. Discriminability analysis

The matched filters described above were used to perform classification for the six mines. The solution to this problem was effected as a bank of matched filters. The classification task was performed by assuming that the most likely target corresponds to the object associated with the filter with the maximum output. Classification was performed in free-space using the matched filters obtained in free space. For mines buried in soil, filters were used based on free air and soil measurments. The matched filter solution assumes that the target signals are known exactly; any uncertainty in the placement of the targets in the environment would substantially degrade

classification performance using this approach. Reference [4] discusses this problem and suggests an alternate approach that incorporates such uncertainty into the classification algorithm.

Table IV lists the discriminability of the various mines. These measures were calculated using a matched filter bank. Half of the target signatures were used to create the matched filter, the other half were used for the analysis. Clearly, it is possible to discriminate each type of mine from the others in free-space or in the ground when the signatures are known exactly. It is not possible to use the free-space signatures as discriminators for the ground measurements. In this case, the classifier can discriminate the Valmara, VS50, and TS50, but the M14, M19 and VS2.2 are always mis-classified as the TS50. It is also important to note that these data were obtained with little to no change in the environmental or geometric variables and thus are not necessarily indicative of real-world performance.

Condition	Valmara	VS50	TS50	M14	M19	VS2.2
Air-Air	100%	100%	92%	100%	100%	100%
Soil-Soil	100%	100%	100%	100%	100%	100%
Air-Soil	100%	100%	100%	0%	0%	0%

Table IV. Discrimination performance of matched filter classifier for signatures estimated in free-space and applied in free-space (Air-Air), signatures estimated in soil and applied in soil (Soil-Soil), and signatures estimated in free-space and applied in soil (Air-Soil).

IV. Conclusions

The results of this experiment indicate that:
(1) The presence of soil does affect the signature recorded by the GEM-3 for low-metallic content mines.
(2) It is possible to detect low-metal land mines using the GEM-3.
(3) The measurements obtained are statistical in nature, not deterministic. Detectors, or algorithms, which effectively incorporate the stochastic character of the signals should be able to out-perform traditional detection algorithms.
(4) The sensor experiences some drift in its response. Drift must be considered both in gathering data and in the analysis. At a minimum, background measurements must be made during data collection for accurate background correction.

V. Acknowledgements

VI. References

[1] *Hand Held Metallic Mine Detector Performance Baselining Collection Plan*, JUXOCO, Fort Belvoir, VA, May, 1998.

[2] Won, I.J., D.A. Keiswetter, and D.R. Hansen, *GEM-3: A Monostatic Broadband Electromagnetic Induction Sensor.* J. Envir. Engin. Geophysics, 1997. 2: p. 53-64.

[3] Won, I.J. D.A. Keiswetter, G. R. A. Fields, and L. C. Sutton, "GEM-2: A New Multifrequency Electromagnetic Sensor", JEEG, Volume 1, Issue 2, August 1996, p. 129-137.

[4] Gao, P. and Collins, L., *"Improved signal processing approaches for landmine detection"*, Detection and Remediation Technologies for Mines and Minelike Targets III Conference, 1998 International Symposium on Aerospace/Defense Sensing and Controls, in press, Orlando, Florida, April, 1998, in press.